EDN®

VOICE OF THE ENGINEER

EM SIMULATION FOR EMC:
KEEPING A LID ON INTERFERENCE
Page 26

ESL SYNTHESIS:
TIPS FOR IMPLEMENTING A VIABLE ESL-SYNTHESIS FLOW
Page 21

DESIGN QUALITY
AND ITS IMPACT ON DESIGN CLOSURE
Page 33

Data acquisition just got a lot easier.

More ways to control. More ways to connect.

The new Agilent 34972A Data Acquisition Switch Unit takes our best-selling Agilent 34970A to the next level. For starters, you get convenient built-in LAN and USB connectivity. Plus, you can control your data acquisition remotely via Web interface. And transfer logged data to your PC with a simple flash drive. No more expensive adapters and connectors. That's easy. That's Agilent.

NEW 34972A	34970A
USB and LAN	GPIB and RS232
Graphical web interface Benchlink data logger software SCPI programming	Benchlink data logger software SCPI programming
3-slot LXI unit with built-in 6 ½ digit DMM	
$1,845 $1,597*	$1,597*

Hurry! For a limited time, get the 34972A at the 34970A price.
www.agilent.com/find/Agilent34972A

Agilent Technologies

Agilent and our Distributor Network
Right Instrument. Right Expertise. Delivered Right Now.

Newark ™
800-463-9275
www.newark.com/agilent

2 GHz Clock Generator

CG635...$2490 (U.S. list)

- **Square wave clocks from DC to 2.05 GHz**
- **Random jitter <1 ps (rms)**
- **80 ps rise and fall times**
- **16-digit frequency resolution**
- **CMOS, LVDS, ECL, PECL, RS-485**
- **Phase adjustment & time modulation**

The CG635 generates clock signals — flawlessly. The clock signals are fast, clean and accurate, and can be set to standard logic levels.

How fast? Frequency to 2.05 GHz with rise and fall times as short as 80 ps.

How clean? Jitter is less than 1 ps and phase noise is better than −90 dBc/Hz (100 Hz offset) at 622.08 MHz.

How accurate? Using the optional rubidium timebase, aging is better than 0.0005 ppm/year, and temperature stability is better than 0.0001 ppm.

You would expect an instrument this good to be expensive, but it isn't. You no longer have to buy an rf synthesizer to generate clock signals. The CG635 does the job better—at a fraction of the cost.

Plot shows complementary clocks and PRBS (opt. 01) outputs at 622.08 Mb/s with LVDS levels. Traces have transition times of 80 ps and jitter less than 1 ps (rms).

Stanford Research Systems

EDN 7.15.10
contents

ESL synthesis: tips for implementing a viable ESL-synthesis flow

22 Here's a checklist of what to do and what not to do when performing ESL synthesis. The items discussed help set realistic expectations for those about to take the plunge.

By Nancy Wu, Gary Smith EDA

EM simulation for EMC: keeping a lid on interference

26 Field-solver simulation can do your work, but you have to do the thinking.

By Paul Rako, Technical Editor

pulse

Dilbert **13**

(a) MISSING ISOLATION CELL IN SCAN PATH — ALWAYS-ON DOMAIN
(b) ISOLATION CELL INSERTED — ALWAYS-ON DOMAIN

Design quality and its impact on design closure

33 Steps to ensure quality early in the design can speed closure and avoid failed silicon.

By Piyush Sancheti and Sanjay Churiwala, Atrenta Inc, and Rob Knoth, Magma Design Automation

DESIGNIDEAS

contents 7.15.10

52

55

DEPARTMENTS & COLUMNS

EDN online contents

www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Path-specific derating to reduce timing pessimism: a proposal
Reducing your pessimism about timing uncertainty as the design progresses can speed timing closure.
→www.edn.com/100715toca

The not-so-smart grid
What exactly are the standards for the next-generation utility grid? In many cases, nobody knows.
→www.edn.com/100715tocb

EUV lithography cannot come soon enough
Sematech's biennial Litho Forum survey shows that the cost of ownership is a leading concern as respondents anticipate their critical-layer lithography needs. Collaboration is key for bringing in viable solutions.
→www.edn.com/100715tocc

What's happening in the electronics world?

The EDN.com News Center provides real-time breaking news and analysis on the global electronics industry, including coverage of semiconductors, lawsuits, IC design, microprocessor units, consumer electronics, analog, and business trends.
→www.edn.com/100715tocd

Exclusive e-book

EDN's third "Designing with LEDs" event, held March 17, 2010, in Santa Clara, CA, addressed high-brightness-LED design challenges from the viewpoint of the hardware engineer. Topics included power control, thermal management, and optics—all of which affect the cost, efficiency, and life span of LEDs. *Six of the technical papers presented at the seminar are available in a downloadable e-book.*
→www.edn.com/100715toce

Our engineers
needed a scope
with digital trigger
to catch every glitch.
So we built one.

See for yourself

www.scope-of-the-art.com/ad/trigger/catchingglitch

BY RICK NELSON, EDITOR-IN-CHIEF

The time is now for 3-D stacked die

As the semiconductor industry moves from "more Moore" to "more than Moore," 3-D-stacked-die implementations will become critical for implementing ever-denser chip packages. Interest in the technology is strong, based on an overflow crowd attending a June 16 DAC (Design Automation Conference) panel, "3-D stacked die—now or in the future?" The consensus seemed to be that the 3-D revolution is imminent. Panelist LC Lu of TSMC (Taiwan Semiconductor Manufacturing Co) said that his company is developing the TSV (through-silicon-via) technology that will interconnect the stacked chips in stacked die, with a focus on design, packaging, and testing—not just the fabrication process.

There are obstacles, however. Panelist Myung-Soo Jang of Samsung suggested that accurate design and analysis tools that work together in a seamlessly integrated flow could speed the adoption of 3-D implementations. Lu agreed that new design methods could help address challenges related to good-die sorting, process variations, and thermal and mechanical stress.

Manufacturers are making progress on the design-tool front. Atrenta, Auto-ESL, Qualcomm, and IMEC at DAC demonstrated a working prototype front-end 3-D chip-design system. The flow the companies demonstrated addresses 3-D-aware high-level synthesis, early partitioning, floorplanning, and multidomain analysis. "The daunting challenges of 3-D design demand a 3-D-aware high-level-synthesis approach," said Atul Sharan, president and chief executive officer of AutoESL.

"Early partitioning, floorplanning, and analysis yield substantial benefits for design predictability on conventional advanced SOCs," said Ravi Varadarajan, Atrenta fellow. "With the emergence of 3-D multitechnology design, this activity now becomes an absolute must-have. You simply cannot hand off a 3-D design to back-end implementation without knowing for certain that it's partitioned correctly."

The demonstration grew out of what Riko Radojcic of Qualcomm called PathFinding technology, which Qualcomm has been developing over a number of years. According to Radojcic, with traditional Moore's law process migration—from 90 to 65 nm, for example—it's relatively easy to project what will happen. The new geometry will yield devices that are smaller, faster, and more prone to leakage—information that can assist in building working and yielding parts.

Such projections aren't necessarily valid or helpful with 3-D parts. Something that allows you to explore knobs at both the architectural and the process ends is necessary, said Radojcic. "That's PathFinding to me." The collaboration with Atrenta, Auto-ESL, and IMEC, he added, is an effort to build a commercial set of tools that assist the PathFinding function.

Pol Marchal, principal scientist for IMEC's 3-D SOC-design initiative, who assisted with the June 14 demonstration, addressed PathFinding technology on June 8 at the IMEC Technology Forum at IMEC headquarters in Leuven, Belgium. IMEC's overall 3-D efforts, he said, involve investigations of TSV technology, wafer thinning and back-side processing, the packaging of 3-D die stacks, cost modeling, and 3-D system exploration, with the last being germane to PathFinding technology.

Marchal called PathFinding a systematic exploration of trade-offs. "Understanding the system requirements provides the engineer with inputs to guide design and technology decisions," he said, adding that a PathFinding flow allows engineers to iterate on design and technology choices to optimize footprint, timing, thermal performance, and other functions.

Marchal cited as an example a mobile consumer device, which would require three or four chip tiers with a package thickness of less than 0.6 mm and more than 1000 TSVs per tier operating at 400 MHz and providing a 12.8-Gbyte/sec data rate, all while consuming less than 2.5 pJ/bit. PathFinding analysis, he said, shows the feasibility of building such a device.

Despite the emergence of 3-D-design tools, obstacles will remain to the widespread adoption of 3-D techniques. DAC panelist Joe Adam of JMA Consulting predicted that incumbents with significant investment in 2-D technologies will be reluctant to change.

Ultimately, however, resistance is futile. As IMEC's Marchal told the June 16 DAC panel attendees, "Practice today or don't play tomorrow." EDN

> ⊞ **Go to the Talkback section at www.edn.com/100715ed to post a comment on this article.**

EDN

CANON **C**OMMUNICATIONS LLC

pulse

INNOVATIONS & INNOVATORS

SMUs support testing of battery-powered products, system components

Agilent Technologies has added two two-quadrant SMUs (source-measurement units) to its N6700 series of modular power systems—the N6781A for battery-drain analysis and the N6782A for functional test. The units, which perform as voltage or current sources or as constant-voltage or constant-current electronic loads, provide the advanced capabilities you need to test battery-powered devices and their components. The company also introduced software and updates for its dc-power-analyzer mainframe.

Thanks to the products' seamless measurement ranging, you can precisely view and measure dynamic-current waveforms without glitches or disruptions. The SMU automatically detects the current of the DUT (device under test) as it changes and switches among its four measurement ranges to the range that returns the most precise measurement.

In combination with the units' built-in 18-bit 200k-sample/sec digitizer, seamless measurement ranging enables 28-bit effective vertical resolution. In one pass and one picture, the units let you see complete current waveforms—from nanoamps to amperes. A complementary feature provides quick glitch-free transitions among source ranges.

When you use the N6781 with the manufacturer's 14585A software, the N6781A provides insights into the battery drain of e-book readers, music players, mobile phones, and pacemakers. The unit also includes an auxiliary DVM (digital voltmeter).

The N6782A can modulate its output to 100 kHz and provides two-quadrant operation, suiting it for such devices as dc/dc converters, power-management units, power ampli-

fiers, and power-management ICs. The fast sourcing and waveform capabilities can stimulate the DUT's input stage, and the electronic-load capabilities can load the output stage and measure its performance, thus enabling comprehensive testing.

The product family includes four mainframes and 24 dc-power modules, providing a range of configurations—from R&D through design validation and manufacturing. The updated N6705B dc-power analyzer has an internal memory of 1 Gbyte and fully supports the new modules, including modules that can output more than 20A.

The N6781A for battery-drain analysis, the N6782A for functional test, and the N6705B mainframe sell for $5300, $4230, and $6900, respectively. A software license for the 14585A costs $1240.**—by Dan Strassberg** ▷**Agilent Technologies Inc**, www.agilent.com/find/n6780.

For use in R&D, the 6705B mainframe houses a variety of modules, including the 6781A and B SMUs, which perform comprehensive tests on battery-powered products and system components.

NI advances IEEE 802.11n WLAN test

National Instruments has introduced Version 2.0 of its WLAN (wireless-local-area-network) software suite for IEEE 802.11n WLAN testing. The suite includes enhanced software tool kits for IEEE 802.11n WLAN-signal generation and analysis. The suite integrates with NI's 6.6-GHz PXIe (Peripheral Component Interconnect Extensions for Instrumentation Express) multichannel RF-signal generators and analyzers to deliver phase-coherent MIMO (multiple-input/multiple-output) RF measurements or IEEE 802.11n WLAN testing. The software also delivers high speed for EVM (error-vector magnitude) and spectrum-mask measurements.

Suite 2.0 adds IEEE 802.11n support to the generation and analysis tool kits to facilitate MIMO testing with true port-to-port phase coherency. For measurements, the tool kit integrates with NI's 6.6-GHz, two-, three-, and four-channel PXIe-5663E VSAs (vector-signal analyzers) to provide accurate multistream MIMO measurements. For signal generation, the tool kit integrates with NI's 6.6-GHz, two-, three-, and four-channel PXIe-5673E VSGs

The new NI WLAN measurement suite adds a variety of software features to improve measurement performance and ease of use.

(vector-signal generators) to generate true phase-coherent, multistream RF signals.

In addition to 802.11n MIMO support, the latest version of the NI WLAN measurement suite improves on the previous version with faster measurement. When you combine it with the latest NI PXIe-8133 quad-core embedded controller, the system can perform composite EVM and power measurements in 7 msec. A result of the tool kit's multicore-enabled measurement algorithms, the fast measurement is crucial for reducing test times and lowering the costs of testing wireless devices.

The suite also adds a variety of software features to improve measurement performance and ease of use. Engineers can software-enable its resolution-bandwidth-filter feature to remove broadband noise and increase EVM-measurement accuracy. With front-end filtering, engineers can achieve residual EVM accuracy of better than –47 dB when testing WLAN devices.

The suite includes soft front panels for signal generation and analysis to support a range of WLAN IEEE 802.11n configurations, including 4×4 MIMO, and to report measurement results and traces. In addition, engineers can use the measurement suite as either a stand-alone executable or an API (application-programming interface) in LabView, in the NI LabWindows/CVI ANSI C environment, or in other C, C++, and .NET development systems.

Prices for the NI WLAN measurement suite for IEEE 802.11a/b/g/n start at $5999, prices for the NI WLAN-generation tool kit for IEEE 802.11a/b/g/ start at $3999, and prices for the WLAN-analysis tool kit for IEEE 802.11a/b/g/n start at $3999.

—by Rick Nelson
▷**National Instruments**, www.ni.com.

IP BLOCK MAKES AUDIO-BANDWIDTH 12-BIT ADC FROM DIGITAL LOGIC

Targeting applications in FPGAs and ASICs, Stellamar has created an IP (intellectual-property) block that uses digital circuitry to make a high-quality ADC. The scheme uses an LVDS (low-voltage-differential-signaling) input as a comparator. By adding a few passive components to the block, you can achieve 12-bit accuracy over a bandwidth of 15 kHz. At 15- and 4-kHz bandwidths, SNRs (signal-to-noise ratios) are 68 and 72 dB, respectively. The design exhibits no missing codes.

One pin of the LVDS pair serves as a single-ended analog input; the other is a feedback pin for the IP block. You can use two of these blocks for a differential input. The analog-input range can sweep through the entire common-mode range of the I/O circuitry, or you can design it to run over a narrower range of voltages with a specified common-mode voltage. The design also works in military applications; you can build it in radiation-hardened CMOS or SOI (silicon-on-insulator). The company will sell the IP block for a fixed NRE (nonrecurring-engineering) fee, plus a per-unit royalty. Stellamar demonstrates the design's operation in many FPGAs.
—by Paul Rako
▶**Stellamar**, www.stellamar.com.

Dilbert By Scott Adams

THE WOODS ARE FULL OF PEOPLE WHO WANT YOUR JOB.

THESE DAYS YOU CAN'T SHAKE A TREE WITHOUT THREE OR FOUR ENGINEERS FALLING OUT.

I'D LOVE TO STAY AND CHAT BUT I NEED TO GO MOTIVATE THE OTHER HEADCOUNTS.

Low-power e-nose checks air quality

IMEC (Interuniversity Microelectronics Centre) and the Holst Centre have developed sensitive integrated sensing elements for gas detection in the form of high-density arrays of polymer-coated microbridges within a MEMS (micro-electromechanical-system) device. The device can detect parts-per-million-level concentrations of vapors using on-chip-integrated read-out techniques. The demonstrated technology is suitable for miniaturization of electronic-nose devices, thanks to its small form factor and power consumption of less than 1 μW/bridge.

Microbridge arrays within a MEMS device constitute an e-nose. A polymer specific to a particular gas coats each microbridge.

Mercedes Crego-Calama, principal researcher at Holst Centre/IMEC and program director for Holst Centre's Human++ program, says that each MEMS bridge within an e-nose device acquires extra mass in the presence of volatile chemical vapors, changing its piezoelectric characteristics. Arrays of such MEMS bridges, she says, can resolve and identify multiple chemical contaminants in an environment and help warn allergy or asthma suffers to avoid potentially harmful environments.

The MEMS-bridge architecture operates at lower power than alternative cantilever-based approaches. Crego-Calama also describes a low-cost manufacturing approach, in which commercial ink-jet-printing technology can deposit the polymers sensitive to specific chemical vapors onto the MEMS structures. The e-nose can monitor air quality, identify pathogens, monitor food ripeness or spoilage, monitor physiological conditions—from breath analysis, for example—and detect biological or chemical weapons.

Within an e-nose device, suspended structures vibrate individually, and changes in their modes of vibration, or resonances, indicate the absorption of a vapor as each bridge device's unique polymer coating determines. Work is ongoing to integrate the structures with low-power analog readout circuits and to demonstrate simultaneous measurements from multiple structures.

—**by Rick Nelson**
▷**IMEC**, www.imec.be.
▷**Holst Centre**, www.holstcentre.com.

TOOL GIVES EARLY AREA, POWER, AND TIMING ESTIMATES

Atrenta has announced the SpyGlass-Physical software tool, which enables RTL (register-transfer-level) engineers to achieve faster design closure by modeling physical-implementation effects at the RTL stage of the design. Previous members of the SpyGlass family provide information on whether a design is syntactically correct and testable; the new member provides early estimates of area, power, timing, and routability for RTL designers who lack physical-design expertise or tools.

According to Ravi Varadarajan, Atrenta fellow, the new tool does not replace a customer's favorite implementation tools but rather makes them more efficient, allowing designers to identify problems earlier in the design cycle.

In addition, the time penalties of successive iterations using traditional tools—in which you find problems after place and route, for instance—limit the amount of exploration designers can do, with the result that they might settle for a suboptimal solution. SpyGlass-Physical alleviates that problem.

François Rémond, director of CAD at STMicroelectronics (www.st.com), describes his use of the SpyGlass-Physical tool in the design of a 55-nm set-top-box chip, which includes 209 million transistors, 230 clocks, 7 million placed instances, 500 signal pads, 53 RTL-IP (intellectual-property) blocks, and 160 hard-IP blocks. He explains that the timing and physical closure of such chips represent a big challenge. "We needed a tool that would partition our SOCs [systems on chips] based on our requirements and provide trade-off analysis and guidance for our implementation tools," Rémond says. "The SpyGlass-Physical product achieved that [goal] on 40- and 32-nm SOCs in significantly shorter time than we expected."

SpyGlass-Physical runtimes can range from two to three hours, which represents a large savings over finding problems deep in the implementation. The product helps to achieve performance targets in concurrent block- and SOC-development processes by using a set of interactive implementation-analysis features.

The result is enhanced guidance for the implementation of both IP blocks and full-chip SOCs. SpyGlass users can easily integrate the product into their design flows and realize the benefits of early physical-implementation modeling.

"Chip-design companies have a great need to reach faster design closure than ... current flows [support]," says Ajoy Bose, PhD, chairman, president, and chief executive officer of Atrenta. "The ability to model the impact of physical implementation on the design at an early stage is a critical aspect that is missing from today's RTL flows."

The SpyGlass-Physical product addresses this gap by providing early estimates of area, power, timing, and routing congestion.—by Rick Nelson
▸**Atrenta**, www.atrenta.com.

07.15.10

Open-source development board: $4.30

Open-source hardware appeals most strongly to developers who want or need to access all parts of their hardware and software platform. To address that requirement, Texas Instruments' new Launchpad development board puts a fully realized 16-bit microcontroller platform within the grasp of any designer for $4.30.

The platform employs TI's Value Line family of low-cost, 16-bit microcontrollers and comes with the 14-pin Value Line

The $4.30, open-source Launchpad development board puts a 16-bit micro-controller platform, including a 10-bit ADC, comparators, and an internal temperature sensor, within the grasp of any designer.

MSP430G2211 and MSP430G2231 processors. Demo firmware on one prepro-grammed processor demonstrates the use of on-chip peripherals. You can swap the processors out in the on-board 20-pin DIP socket as needed. An integrated, USB-powered emulator permits flash program-ming, firmware debugging, and serial com-munication. It can also serve as a program-mer or a debugger for any microcontroller in the MSP430 family. TI offers the free, down-loadable Code Composer Studio and IAR Embedded Workbench KickStart compiler/debugger software packages.

—by Margery Conner
▷**Texas Instruments**, www.ti.com.

pulse

Equivalent-time-sampling scope improves accuracy of characterizing high-speed designs

Agilent Technologies has expanded its DCA (digital-communications-analyzer) portfolio with the introduction of the 86100D DCA-X (extended) wide-bandwidth equivalent-time-sampling oscilloscope. The instrument provides the tools you need to easily and accurately characterize high-speed digital designs. The DCA-X mainframe platform lets you make next-generation measurements that include integrated de-embedding.

The unit is backward-compatible with all previous DCA models and is fully code-compatible with the 86100C mainframe. Future models will support 16 channels, which will enable you to efficiently test parallel designs and will achieve higher throughput in manufacturing.

The DCA series' high analog bandwidth, low noise, and low measurement jitter help you to see the true performance of your designs on signals whose data rates range from 50 Mbps to more than 40 Gbps. You can configure the modular DCA platform for optical, electrical, TDR/TDT (time-domain-reflectometry/time-domain-transmission), and S (scattering)-parameter measurements. The DCA-X improves on the accuracy and ease of use of the 86100C DCA-J (jitter) by integrating de-embedding and embedding using the 86100D-Sim InfiniiSim-DCA software license.

Usability enhancements include graphical-signal processing and dual user interfaces—the FlexDCA, a new, customizable, vector-based interface for scope, eye, and jitter measurements, and the DCA-J. The device also features as many as 64 simultaneous measurements and data-dependent pulse-width-shrinkage, uncorrelated-jitter, and pre- and de-emphasis measurements.

Hardware enhancements include vertical-gain and offset controls for all channels and functions, a user-defined multipurpose button, a user-defined analog-control knob, a faster CPU, and support for as many as 16 measurement channels.

As data rates increase, fixtures and probes play a larger role in degrading waveform fidelity and reducing measurement margins. During verification or compliance testing, you must often remove these unwanted effects to see your devices' true performance. It is common to measure a clean transmitter output and to simulate this signal at the output of a typical channel-loss model. The 86100D-Sim InfiniiSim-DCA waveform-transformation tool set, a new 86100D software option, allows you to quickly and accurately perform both of these tasks.

Agilent has also upgraded its N101A FlexDCA remote-access software so that it now supports both the 86100C DCA-J and the 86100D DCA-X mainframes. The N1010A, a PC-based version of the DCA-X user interface, extends control of the scope through a LAN to a remote lab or home-office PC and allows you to analyze waveforms offline.

US prices for the 86100D DCA-X start at $21,500. US prices for the 86100D-Sim InfiniiSim-DCA waveform-transformation software start at $5000.

—by Dan Strassberg
▷**Agilent Technologies**, www.agilent.com.

The 86100D DCA-X equivalent-time-sampling-scope mainframe, which can acquire signals at frequencies as high as 90 GHz, accepts all of the plug-in modules that its predecessors accept. It also accepts new plug-ins that significantly expand its measurement capabilities and lets you choose between two user interfaces.

Seagate takes cautious approach to solid-state drives

Seagate's new operating-system-agnostic Momentus XT solid-state drives present themselves as conventional 3-Gbps SATA (serial-advanced-technology-attachment) hard drives. Magnetic-media capacities are larger than those of the company's previous offerings, befitting the latest PMR (perpendicular-magnetic-recording) areal density capabilities. Flash-memory capacity is 4 Gbytes, and RAM buffers hold 32 Mbytes. The magnetic media's rotational speed is 7200 rpm, versus 5400 rpm in the first-generation approach.

Seagate claims that Momentus XT prices will track those of conventional hard drives to within a capacity-dependent $50 to $90 markup. Prices for 500-, 320-, and 250-Gbyte units are $156, $122, and $113, respectively.

The drives use SLC (single-level-cell) NAND-flash memory, which is more expensive and more reliable than MLC (multi-level-cell) memory, dispensing with the need for wear-leveling, error-detection and -correction, and erase-block-retirement media-management algorithms. System writes bypass flash memory, so performance limitations don't hamper them. Compare prices on 250- to 500-Gbyte solid-state drives, and the price advantage of Momentus XT will soon become evident.**—by Brian Dipert**
▷**Seagate Technology**, www.seagate.com.

Analog Devices: enabling the designs that make a difference in people's lives.

Measurement, processing, and communications solutions for intelligent energy systems. Optimizing the supply and demand for electricity is a goal we all share. At ADI, we provide data converter, amplifier, processor, and wired and wireless communications technologies to create leading-edge solutions for today's energy market. Our products give consumers precise, real-time smart meter information, enabling them to reduce energy consumption; utilities the ability to accurately monitor grid conditions, increasing system efficiency; and energy producers the ability to optimize solar and wind power generation, reducing fossil fuel dependence. For over 40 years, ADI has brought innovation and imagination to critical designs, including those that make our energy-sensitive world a better place. Learn more at *www.analog.com/energy*.

RF Transceiver

DSP

Metering IC

ADI's family of innovative smart meter solutions offers greater control, communication, and monitoring capabilities.

MAKE**ADI**FFERENCE

www.analog.com

ANALOG
DEVICES

BY HOWARD JOHNSON, PhD

Water analogy

Pump A in **Figure 1** forces water at constant pressure, P_1, around a closed loop. As long as Valve B remains open or at least partially open, the water flowing in the pipe attains steady velocity, V_1. The pump pressure; the viscosity of water; and the effective fluid-flow resistance of the pipes, valves, and elbows used in the circuit determine this velocity. The plumbing circuit is analogous to a simple electrical circuit comprising a battery, a resistor, and a connection from the opposite end of the resistor back to the battery, representing the tank. The electrostatic pressure the battery creates corresponds to water pressure at the head of the pump. The circulation of electrical current in the wires corresponds to the circulation of water in the plumbing.

The concept of resistance translates easily between the two circuits. If you define the fluid-flow resistance of a pipe as the pressure difference from end to end divided by its flow, then a long, skinny pipe makes a high resistance. A big, fat pipe, such as the one still gushing oil into the Gulf of Mexico, makes a low resistance. This fluid-flow equivalence works so naturally that most students, years later at the end of their careers, still think of a high resistance as a long, skinny pipe.

One major difference between the movement of water and that of electrons concerns linearity. The fluid-flow resistance of a piece of pipe varies almost quadratically with the flow, as opposed to the exquisitely linear relationship you normally expect when working with an electrical resistance. Nobody brought that fact to my attention when we discussed the water-flow analogy in school, but, if you want to properly size your lawn-sprinkler system's pipes, you'd better check it out.

The lack of linearity in water-flow systems works to the advantage of a student's understanding. It focuses attention on two fundamental principles of circuit analysis. First, you must always maintain the correct relation of pressure drop to flow for every device throughout a circuit. Second, no water enters or exits a closed system. Those two rules are sufficient to develop a rich understanding of circuit behavior.

On the other hand, the simple linear nature of electronic components tempts electrical engineers to attack every problem with linear matrix analysis and Laplace transforms. These powerful tools confuse beginners. Rather than begin with linear analysis, every electrical-engineering student should first complete a thorough study of nonlinear circuits. Devices such as FETs and bipolar transistors are nonlinear and temperature-dependent anyway, so it makes sense to first introduce nonlinear behavior.

Water- and electron-based circuits, both ultimately dependent on the movement of myriad tiny particles, share many similarities beyond simple pressure and flow relations. For example, if you force too much water through a pipe, the required rise in pressure may burst the pipe, after which nothing flows through to the end and you have to mop up a huge mess. In an electrical circuit, an overtaxed resistor overheats and melts, after which no current flows and you must deal with the smell.

At the other end of the signal-amplitude scale, Brownian motion affects sensitive hydraulic systems, such as the diffusion of ions throughout a nerve ending. Thermal noise affects electrical circuits. It's the same effect either way. We are all just pushing particles around closed loops.EDN

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Email him at howie03@sigcon.com.

Figure 1 Water circulates endlessly around this closed loop, much like electrical current.

Fast Time Division Duplex (TDD) Transmission Using an Upconverting Mixer with a High Side Switch – Design Note 480

Vladimir Dvorkin

Introduction

Many wireless infrastructure time division duplex (TDD) transmit applications require fast on/off switching of the transmitter, typically within one to five microseconds. There are several different ways to implement fast Tx on/off switching, including the use of RF switches in the signal path, or on/off switching of the supply voltage for different stages of the transmitter chain. The advantages of the latter method are low cost, very good performance and power saving during the Tx off-time. In particular, a good place to apply supply switching is at the transmit upconverting mixer because this removes both the transmit signal and all other mixing products from the mixer RF output.

The LT®5579 high performance upconverting mixer fits various TDD and Burst Mode® transmitter applications with output frequencies up to 3.8GHz. Fast on/off supply voltage (V_{CC}) switching for the LT5579 is as simple as adding an external high side power supply switch (note that this technique is equally effective for the lower frequency upconverting mixer, LT5578).

High Side V_{CC} Switch for a Burst Mode Transmitter Using the LT5579 Mixer

The high side V_{CC} switch circuit in Figure 1 uses a P-channel MOSFET (IRLML6401) with an $R_{DS(ON)}$ of less than 0.1Ω. An N-channel enhancement mode FET (2N7002), connected from the drain of IRLML6401 to ground, further improves fall time. The 2N7002's $R_{DS(ON)}$ is less than 4Ω, which is sufficient for this application.

The input driver for the high side V_{CC} switch is a high speed CMOS inverter (MC74HC1G04) capable of driving capacitive loads. The IRLML6401 input capacitance is typically 830pF and the 2N7002 input capacitance is

Figure 1. Upconverting Mixer with High Side V_{CC} Switch

under 50pF. For faster rise times, two high speed CMOS drivers can be used in parallel. Likewise, for faster fall times, a different N-channel MOSFET with lower on-resistance can be used.

With the LT5579 supply current of 220mA, the power supply voltage drop across the MOSFET is only 11mV. The response time of the high side V_{CC} switch is shown in Figure 2. Total turn-on time is only 650ns and total turn-off time is 500ns. These measurements were performed using two RF bypass capacitors at the mixer V_{CC} pin (33pF and 270pF). Higher value RF bypass capacitors can be used, which would result in correspondingly slower rise/fall times.

INPUT CONTROL SIGNAL FOR HIGH SIDE V_{CC} SWITCH

HIGH SIDE V_{CC} SWITCH OUTPUT

LT5579 RF OUTPUT

1µs/DIV

DN4VD F02

Figure 2. V_{CC} Turn-On and Turn-Off Waveforms

The LT5579 upconverting mixer circuit shown in Figure 1 was optimized and tested at an RF output frequency of 2140MHz. The RF output envelope in Figure 2 shows a dip about 300ns after the V_{CC} switch turns on, followed by another, smaller dip at about the 500ns point. Both dips represent the mixer's internal feedback circuit reaction to the ramping supply voltage.

LO leakage to the RF output of the LT5579 was measured at –40dBm when V_{CC} is on and –46dBm for V_{CC} off. The LO port of the LT5579 is internally matched and has a return loss of 10dB to 18dB over a frequency range of 1100MHz to 3200MHz.

When the LT5579 mixer is in the off state, the return loss of the LO port is about 3dB to 5dB across the same frequency range of 1100MHz to 3200MHz. It is advisable to use an LO injection VCO with a buffered output for better reverse isolation, and to avoid any VCO pulling while the LO port impedance changes when switching between the on and off states.

Conclusion

LT5579 and LT5578 mixers without an ENABLE pin can be used in TDD applications with external V_{CC} switching. Using only three parts (IRLML6401, 2N7002 and an MC74HC1G04), a high performance high side V_{CC} switch allows turn-on and turn-off in under 1µs.
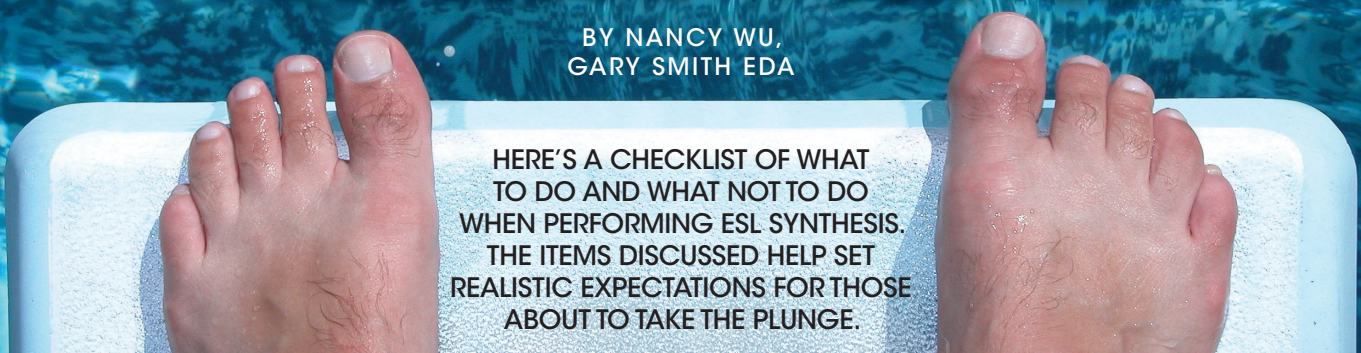
ESL SYNTHESIS:
TIPS FOR IMPLEMENTING A VIABLE ESL-SYNTHESIS FLOW

BY NANCY WU,
GARY SMITH EDA

HERE'S A CHECKLIST OF WHAT TO DO AND WHAT NOT TO DO WHEN PERFORMING ESL SYNTHESIS. THE ITEMS DISCUSSED HELP SET REALISTIC EXPECTATIONS FOR THOSE ABOUT TO TAKE THE PLUNGE.

t would be an understatement to suggest that the ESL (electronic-system-level) market is an enigma. At the 1994 DAC (Design Automation Conference), vendors touted ESL as the panacea that would catapult the next growth spurt in the EDA market. The next 10 to 12 years were a bit disappointing in regard to the mainstream adoption of ESL methods. Depending on how you break up the data, the ESL market in 2005 likely accounted for $150 million in revenues across approximately 20 to 25 companies. Thus, mid-2006 represents a good point at which to assess the ESL user community.

According to early adopters at the time, the top three benefits of ESL were its ability for high-level specification, its strong capabilities for architectural exploration, and the existence of reasonable tools and methods for high-level verification. During the same period, the top three challenges were the fact that ESL offered no clear path to RTL (register-transfer-level) implementation from high-level specification, significant gaps in ESL methodology and flow, and the absence of embedded-software designers from system-level interactions. According to an anonymous user at a major North American electronics-system OEM at the time, the Holy Grail for ESL was an executable and synthesizable specification that designers can simulate and eventually sign off at a high level of abstraction.

By mid-2009, this Holy Grail was still missing. But there has been progress, especially in the synthesizable-specification arena. EDA analyst Gary Smith says that three key segments in ESL are stable and growing: architectural design and exploration, which is primarily the behavioral exploration segment that CoWare and The MathWorks serve; software virtual prototypes, or virtual software development, the tools—from CoWare, Synopsys, Vast, and other vendors—that enable software validation on models of the intended hardware; and ESL synthesis, the tools that generate RTL code from various high-level specifications, including C/C++, and SystemC from companies including Forte, Mentor, and Synfora. Smith calculates that each of these segments grew 30 to 80% on a year-over-year basis during 2008.

As with any other technology, ESL invites squabbling among vendor technologists and early adopters about the goal of synthesis. One goal seems to capture the key aspects from a user perspective: An ESL-synthesis platform enables high-level specification, optimization, and verification and automatically leads to verifiable RTL code with result quality comparable to that of manual methods.

Smith asserts that a viable path to implementation from high-level specification to RTL code enables mainstream users to adopt leading ESL-synthesis tools. In his view, the time for return-on-investment justification has passed; it is now time to educate new adopters about rules of deploying ESL-synthesis tools.

Desaisive Technology Research takes a realistic approach toward measuring and forecasting the ESL market. **Table 1** compares the emerging ESL market with the more traditional RTL market. A migration in revenues from RTL to ESL design is expected during the next three to five years. RTL will not disappear anytime soon. However, the ESL segment should continue to outpace the overall EDA market, and the RTL-design and -verification revenues will likely prove to have peaked in 2007.

Based on input from the technical media and user postings over the past six months, there still appears to be a fervent, vendor-biased debate on one key issue: which high-level language should participants in the EDL market use? This issue is not trivial; it lies at the core of industry participants' objectives for deploying an ESL-synthesis flow.

The choice of high-level language can be a critical factor for successful use of ESL. C/C++ offers the highest level of algorithmic exploration, and designers can easily simulate it with freeware. The greatest challenge is that C/C++

has sequential execution semantics and provides no explicit support for parallel structures. Obtaining good results for the synthesized RTL requires an advanced parallelizing compiler embedded in the synthesis tool to automatically extract parallelism in an application. Thus, the quality of the parallelizing compiler distinguishes one tool from another.

SystemC has emerged as the high-level specification language for hardware modeling and verification. SystemC has seen strong adoption on the hardware specification and, therefore, verification, side in Asia—primarily Japan—and Europe and within a host of consumer-electronics OEMs. SystemVerilog, on

the other hand, is essentially a high-level hardware-verification language, one step higher than RTL. Clearly, the benefits a user hopes to derive from adopting ESL drive the choice of high-level language and, thus, the level of abstraction.

An examination of some state-of-the-art SOC (system-on-chip) architectures may help ascertain the nature and complexity of the challenges facing these design teams. Specifically, consider Texas Instruments' OMAP (Open Multimedia Application Platform) 4 processor, an ARM-based baseband processor for advanced cell phones (**Figure 1**); Broadcom's BCM2153 processor, an ARM-based baseband processor for multimedia applications (**Figure 2**); and STMicroelectronics' STi7167 decoder (**Figure 3**), an advanced set-top box-decoder employing the ST40 CPU/FPU (floating-point unit).
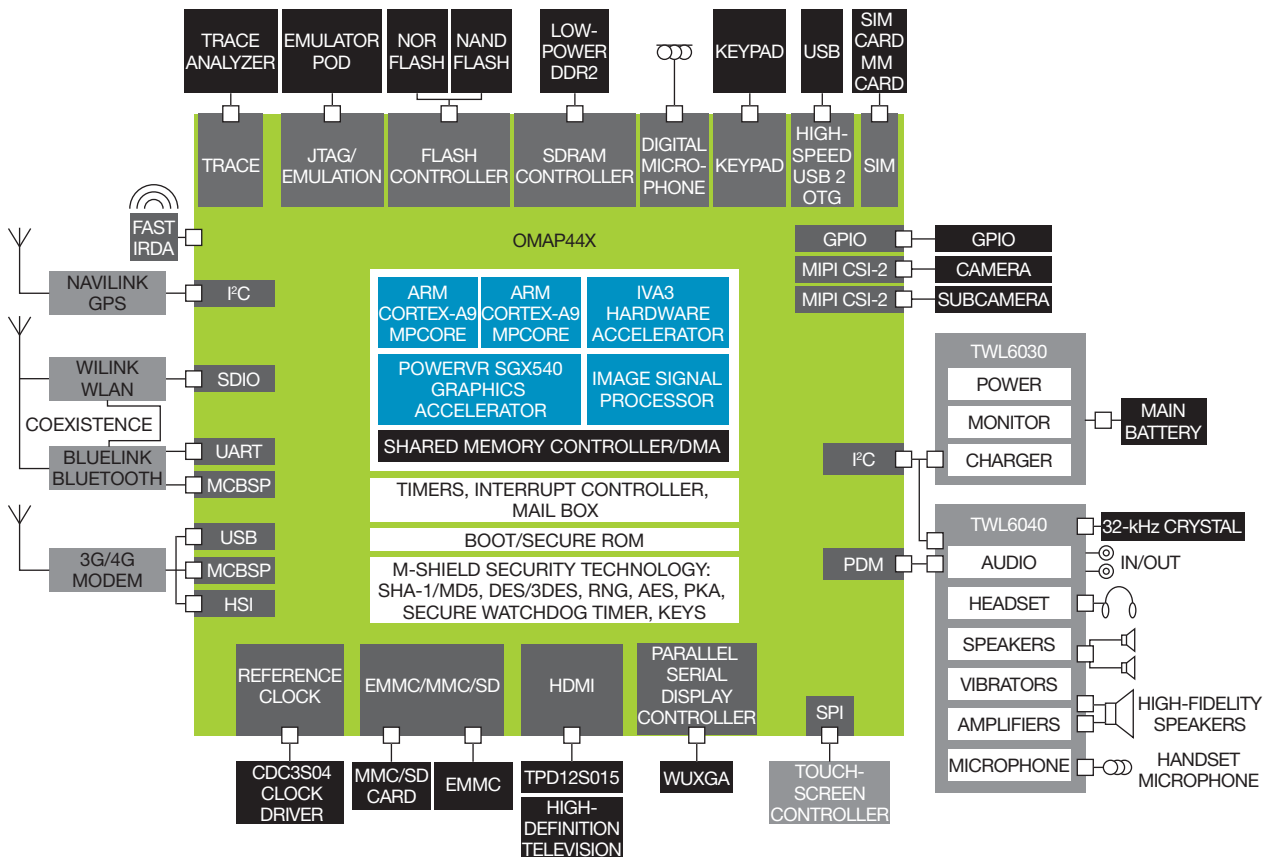
AES=ADVANCED ENCRYPTION STANDARD
DES=DATA ENCRYPTION STANDARD
DRAM=DIRECT RANDOM-ACCESS MEMORY
EMMC=EMBEDDED MULTIMEDIA CARD
4G=FOURTH GENERATION
GPIO=GENERAL-PURPOSE INPUT/OUTPUT
HDMI=HIGH-DEFINITION MULTIMEDIA INTERFACE
HSI=HIGH-SPEED INTERFACE
I²C=INTER-INTEGRATED CIRCUIT
JTAG=JOINT TEST ACTION GROUP
MCBSP=MULTICHANNEL BUFFERED SERIAL PORT
MD5=MESSAGE DIGEST 5
MIPI=MUSIC INDUSTRY PROCESSOR INTERFACE
MM=MULTIMEDIA
MMC=MULTIMEDIA CARD
OTG=ON-THE-GO

PDM=PULSE-DENSITY MODULATED
PKA=PUBLIC-KEY ALGORITHM
RNG=RANDOM-NUMBER GENERATOR
ROM=READ-ONLY MEMORY
SD=SECURE DIGITAL
SDIO=SECURE DIGITAL INPUT/OUTPUT
SDRAM=SYNCHRONOUS DIRECT RANDOM-ACCESS MEMORY
SHA=SECURE HASH ALGORITHM
SIM=SUBSCRIBER-IDENTITY MODULE
SPI=SERIAL-PERIPHERAL INTERFACE
3DES=TRIPLE DATA ENCRYPTION STANDARD
3G=THIRD GENERATION
UART=UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
USB=UNIVERSAL SERIAL BUS
WLAN=WIRELESS LOCAL-AREA NETWORK

Figure 1 TI's OMPA-9 architecture clusters two ARM Cortex A9 processors with key accelerators.

TABLE 1 EDA SEGMENTS REVENUE ANALYSIS ($ MILLIONS)

	2006	2007	2008A	2009E	2010E	5-Yr CAGR
High-level design and verification (ESL)	121.6	139.5	159.5	180	205	
Year-over-year growth	16%	15%	14%	13%	14%	14%
Percentage of total computer-aided engineering segment	5%	6%	7%	8%	9%	
RTL design and verification (traditional)	651.2	696.9	634.6	625	625	
Year-over-year growth	7%	7%	−9%	−2%	0%	1%
Percentage of total computer-aided engineering segment	29%	28%	29%	28%	28%	
Total computer-aided-engineering segment	2213.4	2467	2199.7	2200	2250	
Year-over-year growth	14%	11%	−11%	0%	2%	3%
Total EDA market	4078.3	4464.8	3853.7	3850	3960	
Year-over-year growth	14%	9%	−14%	0%	3%	2%

Source: EDAC MSS Statistics Report and Desaisive Technology Research analysis

All of these architectures have one or two core processors that assume a supervisory role, a main bus structure, embedded memory, and a variety of off-the-shelf IP (intellectual property) for interfacing and control. Each has three to eight application processors or computationally intensive accelerators. These components include hardware accelerators, graphics accelerators, image-signal processors, video accelerators, and security engines. The application processors and accelerators therefore form, in conjunction with the interconnection fabric and embedded-software stack, the basis for competitive differentiation of the SOC. Everything else is essentially off-the-shelf IP.

Such application processors and accelerators, along with interblock communications, dictate what you must specify, verify, and synthesize at ESL. The other leg of system-level differentiation comes from the various stacks of application-specific embedded software. Not surprisingly, programmers debug most such software in C/C++ during the architectural and functional-exploration stage. Thus, true ESL-design specification should begin at the C/C++ level. Proponents of SystemC may disagree with this assessment. However, SystemC primarily finds use in transaction-level modeling that enables performance validation at the system level.

The issue of sequential execution of C/C++ and the challenge of handling parallelism and implicit control also arise. Vendors and leading-edge customers offer several observations.

ANTENNA

RF, POWER AMPLIFIER, AND FRONT-END MODULE

TV OUT

IR

I/O — UART 2

3G DIGITAL RF, ANALOG INTERFACE

BCM2153 HEDGE BASEBAND PROCESSOR

USIM — USIM

SDIO 1 — SD/MMC/T-FLASH/ CE-ATA/MEMORY STICK PRO

ARM11 APPLICATION PROCESSOR

OTP MEMORY

CAMERA — 3-TO-2 MULTIPLEXER / 30-FRAME/SEC QVGA

SDRAM OR SRAM

NAND OR NOR FLASH

EXTERNAL MEMORY CONTROL

HSDPA/WCDMA/EDGE MODEM ACCELERATORS

MODEM AND AUDIO DSP M-STREAM AND SAIC

SECURITY ENGINE

LCD — LCD 1 / LCD 2

GPIO/ KEYPAD — KEYPAD

26-MHz VCXO — PLL

H.264 AND VIDEO ACCELERATORS

JPEG ACCELERATOR

GRAPHICS AND IMAGING ACCELERATORS

AUDIO MIXER FIVE-BAND EQUALIZER

AUDIO — 400-mW STEREO / MICROPHONE / 100-mW STEREO

GPS — SPI

SDIO 2

BCM4328 WLAN

UART 1

PCM

ETM

I²S

BSC

USB 2.0 OTG — 480 MBPS

BCM2048 BT+FM

BATTERY | CHARGER

VIBRATOR | BACKLIGHT

BCM59001 PMU

ATA=ADVANCED TECHNOLOGY ATTACHMENT
BSC=BASE-STATION CONTROLLER
CE=CONSUMER ELECTRONICS
DRAM=DIRECT RANDOM-ACCESS MEMORY
DSP=DIGITAL-SIGNAL PROCESSING
EDGE=ENHANCED DATA RATES FOR GLOBAL EVOLUTION
ETM=EMBEDDED TRACE MACROCELL
GPIO=GENERAL-PURPOSE INPUT/OUTPUT
GPS=GLOBAL POSITIONING SYSTEM
HSDPA=HIGH-SPEED DOWNLINK-PACKET ACCESS
I²S=INTER-INTEGRATED-CIRCUIT SOUND
I/O=INPUT/OUTPUT
IR=INFRARED
JTAG=JOINT TEST ACTION GROUP
LCD=LIQUID-CRYSTAL DISPLAY
OTG=ON-THE-GO

OTP=ONE-TIME PROGRAMMABLE
PCM=PULSE-CODE MODULATION
PLL=PHASE-LOCKED LOOP
PMU=POWER-MANAGEMENT UNIT
SAIC=SINGLE-ANTENNA INTERFERENCE CANCELLATION
SDIO=SECURE DIGITAL INPUT/OUTPUT
SDRAM=SYNCHRONOUS DIRECT RANDOM-ACCESS MEMORY
SRAM=STATIC RANDOM-ACCESS MEMORY
SPI=SERIAL-PERIPHERAL INTERFACE
3G=THIRD GENERATION
UART=UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER
USB=UNIVERSAL SERIAL BUS
USIM=UNIVERAL MOBILE TELECOMMUNICATIONS SYSTEM
SUBSCRIBER-IDENTITY MODULE
WCDMA=WIDEBAND CODE-DIVISION MULTIPLE ACCESS
WLAN=WIRELESS LOCAL-AREA NETWORK

Figure 2 Broadcom's Hedge baseband processor is a single-CPU SOC with a cluster of accelerators.

Most customized logic for new SOCs comprises application processors and accelerators—computationally intensive accelerators with a predominant datapath structure. These architectures put design teams at the forefront of architectural exploration. Architects have for several years been using C/C++ for high-level modeling and specification. The ability to do both high-level specification and functional verification, with an automated path to implementation, brings out the ultimate benefit of adopting ESL. In addition, both Mentor and Synfora have now generated dozens of customers that are willing references that these tools more than adequately handle a flow to RTL. For example, Synfora boasts a unique parallelizing compiler that seamlessly handles both the datapath and the control-logic components of application accelerators.

You must separate the issues of timing verification and performance verification. The focus of high-level verification is functional behavior and the possibility of running more test vectors. Although C/C++ lacks explicit timing information, designers can specify clock frequency and overall performance requirements and can verify them starting with a high-level specification in C/C++. The design team should, however, relegate the issue of detailed timing analysis to SystemC and RTL.

If the goal of adopting an ESL methodology is to begin with the highest level of abstraction with a proven path to implementation, it makes sense to begin at the C/C++ level. Designers can now generate both an RTL model and a SystemC transaction-level model from this specification. You can successfully implement an ESL synthesis flow using C/ C++, including Mentor's CatapultC and Synfora's Pico. For the more than 75% of the design community still using hand-crafted RTL for design specification and verification, the issue is to outline the options available and set the right set of expectations for successful implementation of an ESL methodology.

Mentor and Synfora provide a list of top dos and don'ts that their users have adopted after going through the learning curve of implementing an ESL-synthesis flow. Some common themes advise designers to design at the highest level of abstraction; build architecture awareness into the code at the highest possible level of abstraction; and optimize architecture, system, and power constraints at the high-level design phase. Results matter, and the length of the learning curve and ease of adoption are important productivity considerations. However, you cannot

RESETS
CLOCKS
MODES

ANALOG OUT
PCM I/O
SPDIF OUT

STB PERIPHERALS I/O
EXTERNAL INTERRUPTS

JTAG

CPU/FPU
450-MHz ST40
MMU
INTERRUPTS
INSTRUCTION CACHE
DATA CACHE
TIMER
DEBUGGER

CLOCK GEN
SYSTEM SERV

DUAL FDMA

DISPLAY
BLITTER

CPU/FPU
PCM PLAYERS
PCM READER
DACs

AUDIO DECODER
ST231 CPU

PERIPHERALS
SSC
UART
MAFE
PWM
INFRARED
SMARTCARD

SYSTEM INTERFACES
EMI
PCI
SPI
LMI

FLASH, STATIC FLASH, NOR, NAND

PCI PERIPHERALS

SERIAL FLASH

DDR1/DDR2 MEMORY

STBUS

TRANSPORT SECURITY
DUAL PTI
TRANSPORT-STREAMS MERGER

VIDEO DECODER
ST231 CPU
DELTA MU

DEMODULATOR
DVB-T
DVB-C

DISPLAY
VIDEO PLANES
GRAPHICS PLANES
CURSOR PLANE
MAIN VDP
AUXILIARY VDP
DEINTERLACER
CAPTURE

VIDEO I/O
DVOs
DENC
TELETEXT
DACs
HDMI
VTGs
DVP

CONNECTIVITY

USB

E-SATA

ETHERNET

DAA

TRANSPORT STREAMS
PARALLEL/SERIAL IN/OUT

TRANSPORT STREAMS
INTERFACE IN

DIGITAL VIDEO HD/SD IN/OUT
ANALOG VIDEO HD/SC OUT

CPU=CENTRAL-PROCESSING UNIT
DAA=DATA-ACCESS ARRANGEMENT
DAC=DIGITAL-TO-ANALOG CONVERTER
DDR=DOUBLE DATA RATE
DVB=DIGITAL-VIDEO BROADCASTING
EMI=EXTERNAL-MEMORY INTERFACE
ESATA=SERIAL ADVANCED TECHNOLOGY ATTACHMENT
FDMA=FREQUENCY-DIVISION MULTIPLE ACCESS
FPU=FLOATING-POINT UNIT
HD=HIGH DEFINITION
I/O=INPUT/OUTPUT
JTAG=JOINT TEST ACTION GROUP

LMI=LOCAL-MANAGEMENT INTERFACE
MAFE=MODEM ANALOG FRONT END
MMU=MEMORY-MANAGEMENT UNIT
PCI=PERIPERHAL COMPONENT INTERCONNECT
PCM=PULSE-CODE MODULATION
PTI=PACKET-TYPE IDENTIFIER
PWM=PULSE-WIDTH MODULATION
SD=STANDARD DEFINITION
SPDIF=SONY/PHILIPS DIGITAL-INTERFACE FORMAT
SPI=SERIAL-PERIPHERAL INTERFACE
SSC=SPREAD-SPECTRUM CLOCKING
STB=SET-TOP BOX
UART=UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER

Figure 3 This STMicro set-top box is a single-CPU, multiple-accelerator architecture, but the vendor is vague about the bus architecture.

compromise by accepting results that are not in the same quality range as hand-crafted RTL. Keep in mind that these are company- and application-specific metrics that are difficult to generalize.

Consider verification early and often and how it fits into the overall picture of your high-level-synthesis flow. A key benefit of high-level synthesis is to do most of the verification and debugging at the highest level of abstraction. The vendor's ESL-synthesis tool must thus provide capabilities to debug both functional and performance issues rather than force you to debug at RTL. You should also determine whether the ESL tool generates both RTL and SystemC models, so that hardware designers can perform transac-tion-level-modeling analysis.

Target larger blocks of logic. Other-wise, you derive no productivity bene-fits, and the flow and methodology will not scale to larger blocks or real-life de-signs. You should also use true hierarchi-cal abstraction and ensure that the ESL-synthesis tool supports a multilevel hier-archical method. If the tool internally flattens the design before synthesis, this flattening adversely affects your ability to reuse or retune the design. This loss also affects the issue of results.

Your tools should also provide visibil-ity into what you are building and the downstream design implications. You should have high-level control over ar-chitectural trade-offs and the ability to drive RTL implementation.

Systems comprise datapaths, control logic, and IP. The productivity boost of designing at a high level comes from the deployment of complex systems. Pay close attention to how the ESL-syn-thesis tool handles and optimizes con-trol logic. The synthesis tool should be able to automatically connect the sub-blocks without manual intervention at the RTL.**EDN**

AUTHOR'S BIOGRAPHY
Nancy Wu is chief analyst at Gary Smith EDA. She has a master's degree in econom-ics from the University of Michigan—Ann Arbor and a bachelor's degree in econom-ics from the National Taiwan University.

EM SIMULATION FOR EMC:

KEEPING A LID ON INTERFERENCE

FIELD-SOLVER SIMULATION CAN DO YOUR WORK, BUT YOU HAVE TO DO THE THINKING.

BY PAUL RAKO • TECHNICAL EDITOR

Simulating your product's EM (electromagnetic) radiation will help ensure that you pass FCC (Federal Communications Commission) and CE (Conformité Européenne) tests and will keep your project on schedule. Every product must have EMC (electromagnetic-compatibility) tests. The FCC requires that you test your products to ensure that EM radiation will not cause interference with radios, phones, and TVs. In addition to testing for EM radiation, your product must also exhibit electromagnetic immunity, meaning that a strike from a defined EM pulse will not significantly disturb the product's performance (**Reference 1**).

You need sophisticated software tools to perform EM simulations. These simulations must take into account both small and large features over a broad frequency range (**Figure 1**). You must also select an appropriate simulation method, which can be either a time-domain technique, such as FEM (finite-element method), or a frequency-based one, such as MOM (method of moments). For the largest problems, you need to break the simulations into subdomains or use asymptotic-solutions techniques.

Once you have a powerful computer and the right software, you must place physical and electrical data into the software using database importation or by feeding in mechanical configurations with Gerber and DXF (Drawing Exchange Format) files and manually entering dielectric constants and board-stackup specifications. Finally, you must provide a stimulus to the software, either with Spice or S-parameter data or with a near-field-simulation result from a previous simulation on a subsystem in the product.

SPICE VERSUS FIELD SOLVER

You cannot use Spice to simulate EMC because Spice is a matrix-math computational solver for Kirchhoff's equations that uses lumped-element models of discrete components. At best, you can use Spice to model a lossy transmission line to define what happens to the signal, but it does not reveal which fields radiate into space. For this problem, you need a field-solver simula-

tion (**Reference 2**). A field solver uses finite elements, meshing, and iteration to solve Maxwell's equations for your circuit design. EM-simulation software must account for the mechanical configuration and the materials you use in the design (see **sidebar** "Computer power").

The highest frequencies you are trying to simulate and the size of the circuit dictate the scope of the field-solver problems you will encounter. Wavelengths are 30m at 10 MHz, meaning that a 1-cm trace is much smaller than the wavelength (**Figure 2**). The software would not have to mesh the trace into smaller sections to iterate toward a solution. The 30m wave acts almost uniformly on the 1-cm trace.

Imagine a 10-GHz radar signal with a 2-cm wavelength bouncing off a battleship. The field-solver software must

Gain_Tot[dB]
10.0
6.0
2.0
-2.0
-6.0
-10.0
-14.0
-18.0
-22.0
-26.0
-30.0
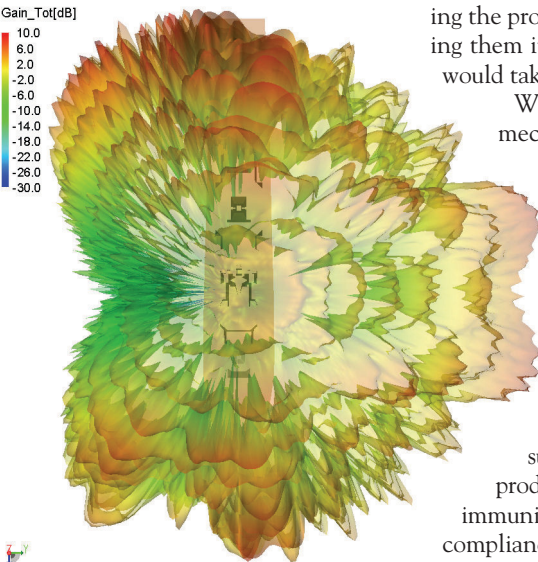
Figure 1 The EM pattern from a small antenna on a large ship exhibits the multiscale problem of an EM simulation (courtesy Feko).

break the battleship into billions of tiny meshes, fitting 10 or even 100 into each square centimeter of the ship's surface. The surface of a metal battleship is not purely reflective, so the software must do 3-D meshing and has even more elements to compute because it must also do the interior areas. The workstation that runs the software needs hundreds of gigabytes of memory to store intermediate calculations for the meshes, and it would take months to solve for

the fields over this large area. You can solve the memory problem by breaking the problem into domains and solving them iteratively, but that approach would take even longer.

When you test for EMC, small mechanical features result in big changes in performance. A slot in a cover, a misrouted trace, or an aluminum heat sink on an IC package can all cause your product to fail EMC-radiation testing. These mechanical features serve as antennas, so they also receive energy from their surroundings, giving your product poor electromagnetic immunity. The standards require compliance to frequencies of 960 MHz and beyond. For this reason, simulating for EMC is a broadband problem with heavy computational requirements. You must simulate for those frequencies; thus, simulating a large system takes an unacceptably long time. The

complexity of the problem is monstrous even for a rather simple product. Also, multiple phenomena, including electrical fields radiating from traces, magnetic fields from inductors, and both types of fields radiating into and from cables, are responsible for EMI (electromagnetic interference).

A typical EM-simulation strategy divides the problem into pieces and depends on both relative and absolute measurements. You need to know how customers will use the product, divide your EMC analysis into manageable pieces, and then evaluate those pieces as they relate to the whole problem. The principle of superposition can be a big help. It states that, for all linear systems, the net response at a given place and time that two or more stimuli cause is the sum of the responses that each stimulus individually would have caused. If three main contributors are affecting your EMI signature, you can individually simulate each one, with different techniques if necessary, and then add the results in an rms (root-mean-square) fashion if they are not related. Sometimes, though, one system affects the other, and they do interact.

Once you have simulated the PCB (printed-circuit board), you represent that simulation as a radiating model that you then plug into a larger assembly. Even if you can use likely signals to simulate the radiation from your PCB, you may also have a few switching power supplies that have not only electric fields, but also magnetic ones. A case surrounds these components, and the cables from the product are antennas that radiate energy to make you fail EMC testing and receive energy to make your circuit fail immunity tests. You may also have to decide how disparate radiation patterns add up to a total emissions level. That decision may bring up the ugly reality of nonlinear

TABLE 1 FINITE-ELEMENT AND TIME-DOMAIN TECHNIQUES

Frequency	Electrical length (wavelengths)	FEM	TD	MOM/ MLFMM	Asymptotic
1 MHz	0.5	X	X	X	
10 Mz	5	X	X	X	
100 MHz	50		X	X	
1 GHz	500			X	X
10 GHz	5000				X

Source: CST

circuits, such as RF-power amplifiers that you drive into saturation to get good efficiency. Superposition techniques don't work in nonlinear systems and may cause you to underestimate the radiation from the circuitry.

SELECTING A TECHNIQUE

The mathematicians and software wizards who work at field-solver companies have developed many methods to help you do EM simulations. You can use 2-D simulation programs, such as HyperLynx and SIwave (signal-integrity wave) to evaluate the EMC of a PCB. Fixing the signal- or power-integrity problems on the card often fixes your EMC problem, as well. You can use time-domain simulations for lower frequencies and smaller physical problems. The key benefit of the time-domain techniques is that they use GPU (graphics-processing-unit) cards, which speed up the math.

James Stack, training, applications, and consulting manager at Remcom Technology Solutions, reports that adding one GPU speeds solvers by a factor of 30 and stuffing your computer with four GPU cards can speed things up by a factor of 150. David Johns, vice president of technical support and engineering at CST (Computer Simulation Technology), reports that his company's time-domain solver runs problems 12 times faster with a GPU.

Unfortunately, at higher frequencies, time-domain techniques are not the best way to solve for EM fields. FEM and time-domain field-solver techniques work best for slower signals,

COMPUTER POWER

As you evaluate software packages to help you with EMC (electromagnetic compatibility), be sure to ascertain which formats the packages import as physical models and which PCB (printed-circuit-board) packages they work with. Ask the vendors how the software represents stimulus functions and make sure that you can create those representations. Finally, buy the proper type of workstation. It should have multiple GPUs (graphics-processing units) if you have the time and the budget for time-domain simulations. If you are doing faster frequency-based solving, then you need multiple cores and multiple CPUs along with a large memory to store the intermediate results as the simulation progresses.

Modern computers can often solve—in hours—complex products, such as computer-sized systems. A key spec is how large a problem the computer can solve overnight—a likely scenario. You spend the day setting up and designing the configuration, and you then launch the field solver before you go home. The next morning, you will be able to look at the solutions and decide on the next iteration of design changes and simulations. Just remember that working on EMC problems is no place for departmental silos. The fixes will be both mechanical and electrical, and the team working on EMC needs the authority to change anything that will solve the problem.

while MOM and asymptotic solvers work for faster speeds and larger problems (**Table 1**). You are better off using a frequency-based solver in a PC workstation with lots of memory and multiple CPU cores. Companies such as Feko and CST use MLFMM (multilevel-fast-multipole-method) techniques, which solve large problems with less computer power. As the problems become large and must run at frequencies greater than 10 GHz, you must use special solvers that can do asymptotic analysis, which solves for large sets. In some simulations, one physical domain affects another (see **sidebar** "Multiphys-

ics keeps tabs on your design").

Some products, such as those from Cadence, Mentor Graphics, and Zuken, have tools to get electrical and physical information into the simulation software. When you do your PCB design in these vendors' tools, the vendors provide a complete representation of the PCB-layer stackup and material, allowing their signal-integrity and field-solver tools to use this data in their simulations.

Also make sure that point tools can accept your PCB data. CST and Sigrity take databases from Cadence, Mentor, Zuken, and Altium, and these tools and many others accept ODB (open-database)++ PCB fabrication to define the physical configuration and materials in PCBs. Full-wave-simulation vendors, such as SPEAG and 2Comu, are familiar with 3-D databases and can import STEP (Standard for the Exchange of Product Model Data), IGES (Initial Graphics Exchange Specification), DXF, and other mechanical-solid-modeling outputs (**Figure 3**). Once in the simulation program, the program meshes the solids with algorithms appropriate to the method.

Defining the mechanical shapes and dielectric constants of the physical design is only part of the EMC-analysis problem. When using a time-domain technique, you can just put the proper

Figure 2 This battleship is meshed for a 10-MHz analysis. A 10-GHz analysis would have a 100-times-finer mesh (courtesy CST).

Figure 3 You can import physical components from mechanical CAD packages and use the solid models in EM simulations (courtesy SPEAG).

time-domain waveforms on the ends of the traces. IBIS (I/O-buffer-information specification), a time-domain look-up table of driver-pin waveforms, can describe the rise and fall of a signal on a pin. You also must define the data on the pin; a PBRS (pseudorandom binary sequence) is often adequate for representing the spectral content of the signal on a functional product.

You can use IBIS-AMI (algorithmic-modeling interface) to define the pre-emphasis circuits and equalizers in the chips you are using, but it does not define the actual waveforms that will appear when your product is running. Typically, you just use a PBRS into the IBIS-AMI blocks. Meanwhile, your design may have hundreds of traces that might interact (**Figure 4**).

S parameters are the best way to represent the spectral content for the sources of EM noise at high frequencies. An S-parameter representation of a PCB block still does not give you the spectral content that the block will output unless you properly excite the block with signals typical of those that the product will use. Using EM simulations for EMC does give rise to a "chicken-or-egg" problem. Sometimes, the only thing that has the adequate representation of the frequency spectrum being radiated is a working board inside a real case. In that situation, it may make no sense to simulate the problem when you can simply test it, but doing simulations is important. You must know where your simulations deviate from actual results, and doing a correlation between the simulations and real measurements allows you to improve your models, your meshing, or your technique. This approach may not save time for the product you are currently working on, but it can shave months or even years off the development of the next one. Getting results from a Spice simulation that creates a near-field model that you then import to a 3-D solver is a good way to stay in control of your EMC problems.

TOOLS FOR EM SIMULATION

No single piece of software can do EMC analysis. You should assemble a suite of software tools to help battle your EMC demons. For example, board-level tools can ensure that the signals go where you intend instead of radiating to space. All enterprise-class PCB companies have good field-solver tools to help with signal integrity (**Reference 3**). Mentor Graphics may be most well-known for its HyperLynx tool, but Cadence, Ansoft, and Zuken also have powerful tools that work on a PCB with hundreds or thousands of traces. SiSoft makes a signal-integrity tool similar to HyperLynx. Sigrity Systems offers its software as a point tool to plug into PCB flows. This tool finds how power-integrity problems and signal-integrity problems relate to each other (see **sidebar** "SI and power integrity are also important"). Once you have a well-designed PCB, you may then have to approach the problem as if you were an RF-board designer.

RF-design-software companies Agilent ADS, AWR Microwave Office, Ansoft, Sonnet, CST, and dozens of others can help you deal with the vagaries of EMC analysis. Most of these companies also offer plug-in software, such as the EMPro software in Agilent's ADS, which performs EM modeling. These tools also account for metal boxes and shields around the circuits and can evaluate the relationship between the electrical and mechanical aspects of your design—an inherent requirement of RF design. RF designers know that their circuits' performance changes after the cover is on. RF-design tools can model the cooling slots in the case and tell you the amount of radiation coming from them for a given frequency excitation.

Excitation constituting random fields pouring from your PCB as it operates is a more intractable problem, but field-solver companies CST and Ansoft demonstrate how you can solve

Figure 4 Crosstalk from the signals on the bottom traces of a PCB creates surface currents on the upper left (courtesy CST).

MULTIPHYSICS KEEPS TABS ON YOUR DESIGN

With experience and deliberation, you can break up your EMC (electromagnetic-compatibility) analysis and use software at all points in the process to ensure you pass FCC (Federal Communications Commission) and CE (Conformité Européenne) testing. Remember, though, that surprises can always occur. The term "multiphysics," relating to the interactions of domains in simulation, often comes up in EMC. For example, if the RF energy heats a material that exhibits a Curie point, then the simulation must account not only for the electrical behavior of the space around that material, but also the thermal effects on the material. It must then calculate the change in fields that the changes in the magnetic properties of the material cause.

A more common multiphysics problem occurs when slots in a product's case affect both the thermal behavior of the product and the amount of RF that exits the case. In these cases, mechanical engineers and electrical engineers, respectively, debate using large openings in the case to decrease heat versus using no openings to attenuate EMI (electromagnetic interference). Multiphysics-capable software from Ansys and Comsol, for example, can mediate these conflicting departmental desires. Letting the software simulate an optimum point for cooling that also provides some margin for EMC performance at least stops the bickering and lets the product design move forward. Again, the software can't guarantee that you will pass FCC testing, but it can provide a compromise from which both groups can evaluate further work.

it. You use time-domain simulations with real waveforms on multiple traces to do a simulation. You then capture a near-field representation of the radiation from the PCB. At a short distance from a source, the electric and magnetic fields do not directly relate, as they do in a wave propagating through space. You then plug this near-field result into a full-wave solver that can calculate the effects of your product's case, cables, and other mechanical features (**Figure 5**).

SOFTWARE CAN'T THINK

Field-solver EM simulations do not provide an EMC panacea. Furthermore, they are not magical genies that can solve a design disaster. Field-solver vendors stress that computer simulations are parts of the entire design process, not some tacked-on afterthoughts that you do as a penance when your product fails FCC testing. You can't expect a computer simulation to identify every area in which you may have to make improvements. However, if you use and understand the simulations of various parts of your design as it progresses, you will be in better shape when submitting your product for FCC and CE testing. In many cases, the most valuable thing that a computer simulation will do is teach you the nonintuitive behavior of EM fields in a complex product. Playing with the configurations, materials, and shielding will help you understand what is going on, and you can design the product to comply with regulations.

With signal frequencies in the gigahertz, a finned heat sink on an FPGA acts as a phased-array antenna, radiating energy in your product. The cooling slots on the case are also phased antennas. Even if you do not have the time or budget to do a full simulation of the electrical signals on the board radiating to a point 3m away, you can still use full-wave simulations. A broadband simula-


Figure 5 Use a 2-D solver to calculate the near fields from a PCB. You can then use the near-field data as a source in a 3-D solver to yield EMC-performance data outside an enclosure (courtesy Ansoft).
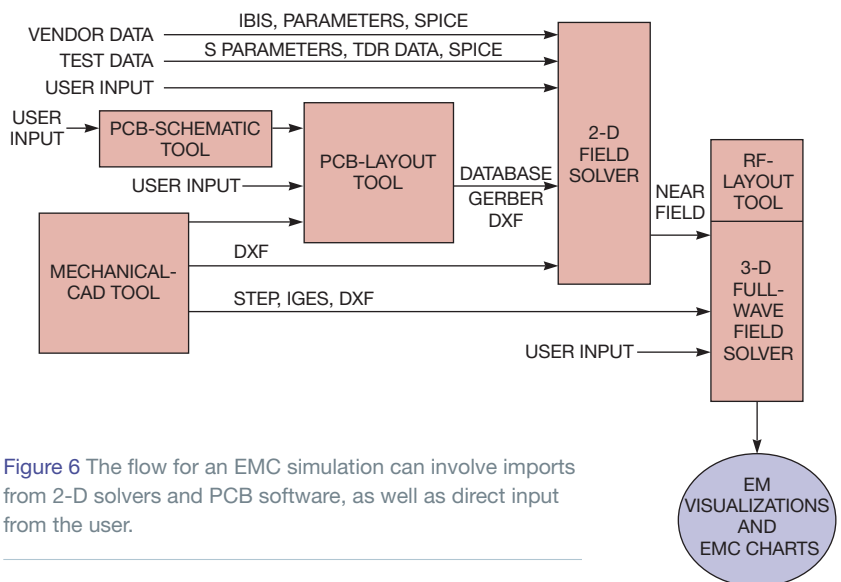

Figure 6 The flow for an EMC simulation can involve imports from 2-D solvers and PCB software, as well as direct input from the user.

tion of the heat sink tells you at which frequencies the sink resonates and the spatial pattern of the resonance. You can also do a broadband excitation of the slotted case. If the frequencies and locations of the heat-sink resonance align with the resonance of the case, those frequencies will cause trouble. The fix may be as simple as rotating the heat sink 90° or changing the spacing of

SI AND POWER INTEGRITY ARE ALSO IMPORTANT

Field solvers are useful for evaluating the SI (signal integrity) of your circuits. Solvers can also help with the related problem of power integrity on your PCBs (printed-circuit boards). These problems are related to EMC (electromagnetic compatibility) because solving SI and power-integrity problems often solves radiation and susceptibility problems, as well.

Lawrence Williams, director of product-management groups at Ansoft, says that SI, power integrity, and EMC reside under a common "product-integrity" umbrella. You can often use a simpler 2-D solver to help you design a PCB with good signal integrity. EMC is a 3-D, full-wave field-solver problem, but, if you optimize the signal and power integrity on the board, it radiates less.

More than a decade ago, Mentor Graphics' HyperLynx tool predicted the performance on FCC (Federal Communications Commission) tests of a complex PCB. Since then, Mentor Graphics has acquired—and is adding to HyperLnyx—3-D field-solver software from Zeeland. This approach will help solve for vias and connector stubs that carry signals that are faster than 6 GHz. At that speed, even PCBs need 3-D solvers. Just as signal integrity relates to EMC, reducing radiation often also reduces immunity problems.

Figure 7 When your EMC simulations correlate with the real world, you can speed up and eliminate risk from product development (courtesy Remcom).

the fins, the slots on the case, or both.

Field-solver programs have steep learning curves, especially for engineers unfamiliar with 3-D simulations (**Figure 6**). Once you understand the software, you must learn how to import your physical configurations and electrical stimuli. It may seem like an unending task, but once you get a simulation to accurately predict the EMC performance of your product, you will see the attraction of using simulations (**Figure 7**). They allow you to evaluate things in hours instead of months. They don't guarantee that your product will pass radiation and immunity tests, but they give you a big head start over companies that simply use "cut-

and-try" methods to get their products through FCC and CE approvals. You do testing at the end of the product cycle, when whether you ship the product determines your company's fortunes.

Smart engineers use software simulations to evaluate EMC in the early phase of the design cycle, moving the risky EMC problems away from the critical product-release phase. They still need to make design and schedule changes but have enough time to efficiently solve the problem without delaying the product's introduction.**EDN**

REFERENCES

1 Rako, Paul, "RFI: keeping noise out of your designs," *EDN*, Jan 10, 2008, pg 25, http://bit.ly/bP6kjb.
2 Rako, Paul, "Beyond Spice: Field-solver software steps in for modeling high-frequency, space-constrained circuits," *EDN,* Jan 18, 2007, pg 41, http://bit.ly/awRxyd.
3 Rako, Paul, "Swimming in the channel," *EDN*, March 18, 2010, pg 34, http://bit.ly/cT7QVB.

FOR MORE INFORMATION

Agilent
www.agilent.com

Altium
www.altium.com

Ansoft
www.ansoft.com

Ansys
www.ansys.com

AWR
www.awrcorp.com

Cadence
www.cadence.com

Comsol
www.comsol.com

CST
www.cst.com

Feko
www.feko.info

Mentor Graphics
www.mentor.com

Remcom
www.remcom.com

Sigrity
www.sigrity.com

SiSoft
www.sisoft.com

Sonnet Software
www.sonnetsoftware.com

SPEAG
www.speag.com

2Comu
www.2comu.com

Zuken
www.zuken.com

You can reach Technical Editor Paul Rako at 1-408-745-1994 and paul.rako@ cancom.com.

BY PIYUSH SANCHETI AND SANJAY CHURIWALA • ATRENTA INC
ROB KNOTH • MAGMA DESIGN AUTOMATION

Design quality and its impact on design closure

STEPS TO ENSURE QUALITY EARLY IN THE DESIGN CAN SPEED CLOSURE AND AVOID FAILED SILICON.

The cost of SOC (system-on-chip) design continues to skyrocket, market windows continue to shrink, and design complexity continues to grow exponentially. These challenges are only a few of those that SOC designers face. In an effort to prevent major disasters, designers must ensure that the SOCs achieve design closure, which includes meeting certain key objectives, such as performance, power, and area. Design-closure objectives are often in conflict with each other, however. Designers must constantly trade off one for the other to ensure that the design stays within the end-user application's requirements.

A typical SOC design starts with an RTL (register-transfer-level) description to capture user intent and a set of design constraints to drive implementation. The design team first verifies the RTL for correctness of functional intent through simulation and formal verification. The design then goes through a series of implementation steps, including synthesis and placement and routing, which eventually result in a GDSII (Graphic Design System 2) layout for silicon manufacturing. The quality of incoming design and the associated constraints have a large impact on the designers' ability to reach closure. However, you can ease this process by using a series of design-quality measures starting at RTL and continuing throughout implementation, focusing on quality measures at the five stages during an integrated RTL-to-GDSII implementation flow (**Figure 1**). You can expand the concept to other stages of implementation or adapt it to other flows, including presynthesis-RTL quality; postsynthesis, postscan-netlist quality; post-timing-netlist quality; post-placement-netlist quality; or postroute-netlist quality.

PRESYNTHESIS-RTL QUALITY

SOC designs that don't start well will usually fail to reach closure. Qual-

ity measures at the RTL stage of design go a long way toward determining successful design closure and working silicon. Once you synthesize the design, you are to a large extent freezing the design intent, and you have limited flexibility for correcting design-quality issues inherent in RTL.

Modern SOC designs typically cater to multiple end markets to amortize the high cost of design. The same design can have multiple variations and live on for multiple generations through updates and upgrades. This scenario is prevalent in consumer electronics and automotive chips, in which manufacturers accomplish 80% or more of the de-
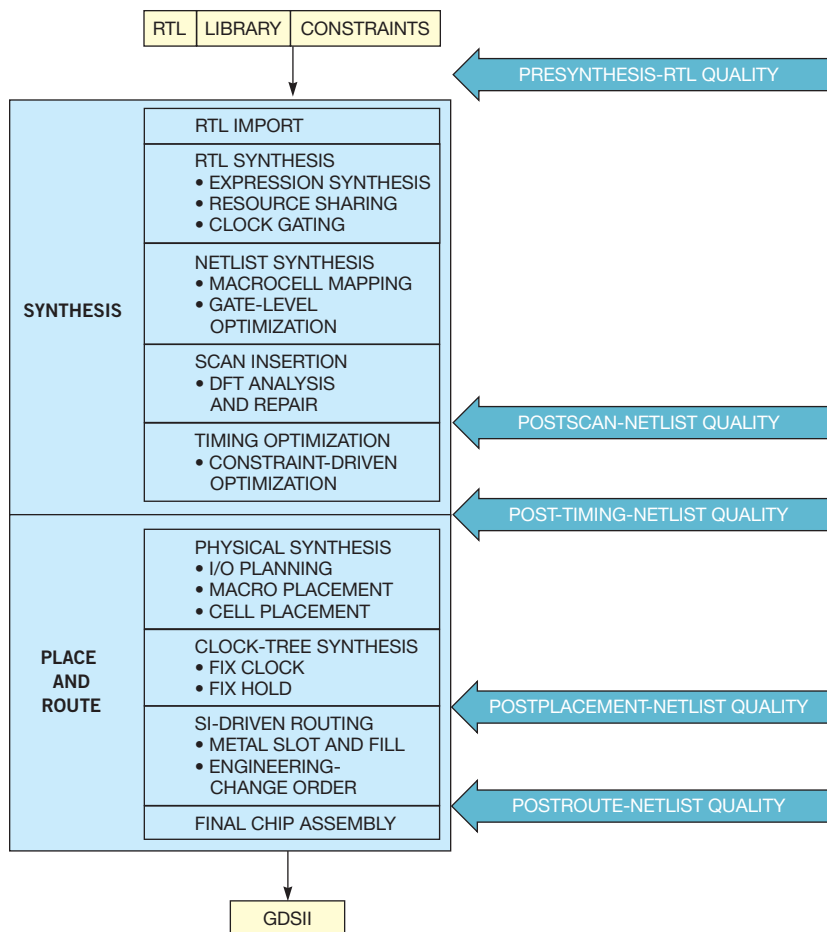

Figure 1 A typical SOC implementation flow should include quality steps.

sign through reuse. Future generations may reuse the RTL you create for the current design and therefore it could have a longer shelf life than the current design. You must also consider commercial third-party IP (intellectual property), such as processors, digital-signal-processing blocks, and bus fabrics, and interface IP, including Ethernet, USB (Universal Serial Bus), and PCI (Peripheral Component Interconnect). The SOC team typically receives this IP in RTL.

For these reasons, you must ensure the quality of the RTL and constraints going into synthesis. Design teams often focus on functional correctness using simulation and formal verification, but spending some effort on implementation feasibility and the overall quality of RTL can go a long way toward accelerating design closure. Design teams can achieve such quality measures through a set of analyses on the RTL and design constraints.

STRUCTURAL AND CONNECTIVITY INTEGRITY

RTL linting can weed out syntax and semantic issues and ensure compliance with coding standards. RTL designers should, however, address more serious structural and connectivity issues at this early stage. If left undetected, they may later lead to more serious design-closure issues. Examples of these issues include excessive levels of logic between flip-flops (**Figure 2**), combinational loops, unintentional latches,

Figure 2 Excessive levels of logic between flip-flops are potential timing-closure problems.

blocking assignment in sequential blocks, variables or nonconstants in the terminating condition of a loop, missing asynchronous resets from a sensitivity list, multiply-driven nets without a tristate, undriven nets and ports, and mismatching between the left- and the right-hand sides of an assignment. Although you may be able to detect and fix some or all of these issues during synthesis or later stages of implementation, it is more efficient to fix them before putting any effort into implementation.

CLOCKS AND RESETS

A typical SOC contains heterogeneous IP from different sources. As a result, the number of asynchronous clock domains on a chip has increased dramatically. It is possible for one chip to have 20 or more clock domains. You must ensure that clocks and resets are properly designed. When data signals cross between asynchronous clock domains, you must synchronize them to prevent metastability (Figure 3). Clock synchronizers can range from multiple-flip-flop synchronizers to more exotic schemes, such as FIFO (first-in/first-out) buffers with handshaking. It is important to prevent the data loss and reconvergence of synchronized signals to ensure reliable behavior. You must synchronize deasserted resets, even if they are asynchronous, with the clock domain.

You should ensure not only that synchronizers are in place on crossings but also that you've correctly implemented the

(a)

(b)
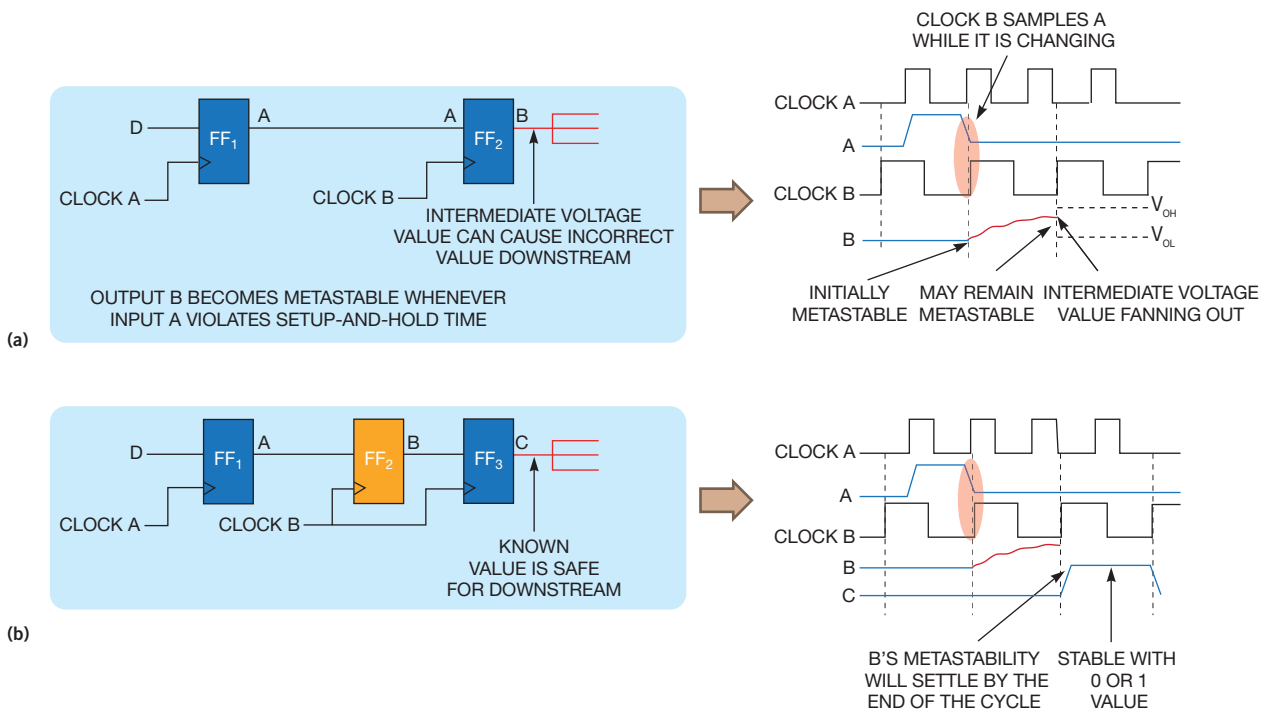
Figure 3 A metastability problem can occur in clock-domain crossings (a). Designers often solve the problem using another approach (b).

Figure 4 Various clock-gating opportunities are available.

additional gating logic can consume more power than the power you save by gating the clock. Excessive clock gating can lead to timing-closure issues and contribute to routing congestion. You should instead selectively apply clock gating in places in which it has the most impact on power.

RTL analysis for clock gating can help in a number of ways. At RTL, you can identify global clock-gating signals, which can gate clocks for an entire design or for large register banks. A review at RTL can also analyze and prioritize explicit clock enables. RTL designers define these enables for their power-saving potential and help remove those that save little or no power. Power-management designers can also identify new or implicit clock-gating opportunities that RTL designers may have overlooked. In addition, power-management specialists can also generate directives for synthesis to intelligently gate clocks.

Various clock-gating opportunities are available to RTL designers (**Figure 4**). Power designers can do similar analysis to identify data-gating opportunities, in which a cloud of combinational logic drives an enabled register. Gating the combinational logic using the same enable that you apply to the terminal register eliminates wasted power resulting from

YOU MUST ADDRESS POWER MANAGEMENT EARLY IN THE DESIGN FLOW, USING ARCHITECTURAL TECHNIQUES AND RTL TECHNIQUES.

toggles in the combinational logic when the register is disabled. For example, an N-bit multiplier, with the input data bits arriving at different times, is a candidate for data gating. The multiplier continues to multiply even though the results remain unused until all the bits of both data inputs have arrived. Data gating can be an effective technique for such datapath-intensive designs that digital-signal processing commonly uses.

DESIGN FOR TEST

Designs must have high test coverage, both for stuck-at- and at-speed-fault modes, especially in consumer electronics, which must quickly reach silicon volume with few defects. Traditionally, design teams stitch scan chains during synthesis or later stages and test coverage and then estimate test coverage using ATPG (automatic-test-pattern-generation) tools. However, most testability issues are detectable and correctable at RTL, so the design will eventually meet test-coverage goals.

For example, the key to high stuck-at-fault coverage is to make sure that the design is fully controllable and observable in scan mode. However, high stuck-at-fault coverage in RTL encounters many barriers, including nonscannable flip-flops whose inputs are unobservable and whose outputs are uncontrollable. Designs that internally generate control signals, such as clock or asynchronous set/clear, are the most common causes of this situation. Nontransparent latches are oth-

protocol. For example, a FIFO should have no overflow or underflow, and you must implement proper sequencing between requests and acknowledgments in a handshaking scheme. Functional simulation may not detect clock-domain-crossing bugs unless verification engineers create dedicated testbench scenarios for each crossing, a daunting task for designs that have thousands of such crossings. You must employ structural-analysis and formal-verification techniques to exhaustively analyze and verify clock-domain crossings.

POWER REDUCTION

Power has come to the forefront of design-closure concerns for a variety of reasons, including battery life, cooling costs, reliability, and energy efficiency. Studies show that the determination of more than 80% of the power of a design happens by the time it enters synthesis. For that reason, you must address power management early in the design flow, using architectural techniques, such as multiple voltage domains, power domains, and dynamic-voltage-frequency scaling, and RTL techniques, such as clock and data gating. Designers must start with an estimate of the power consumption of the design and selectively apply these techniques based on power targets for the design.

Voltage and power domains add new challenges for design closure. In voltage domains, it is important to insert level shifters wherever signals cross from one voltage domain to another. Similarly, you must place isolation cells in power domains that can be shut off when not in use to ensure that unpowered outputs are not floating. These floating signals could introduce a functional error or a high-leakage path to ground. You must also ensure that the enable logic for isolation cells is in an always-on domain. Some designers insert level shifters and isolation logic at RTL, and others capture the power intent in CPF (Common Power Format) or UPF (Unified Power Format) for automatic insertion by downstream implementation tools. In either case, designers must ensure that the design has level shifters and isolation-logic cells at each such crossing.

With judicious use, clock gating can be an effective power-reduction technique. Most synthesis tools can automatically insert gating based on clock enables in the RTL. However, not all clock gates save power, especially in the case of registers, such as flip-flops, that are almost always enabled or if the design has only a few gated registers. In such cases, the

Avago SoC Brings Single Package Bluetooth 2.1 and Laser Tracking to Wireless Mice

Introduction

Avago Technologies' new ADNS-7630 is a fully integrated, feature-rich Bluetooth® (BT2.1) System-on-Chip (SoC) LaserStream™ navigation sensor for wireless mouse applications. It is the first laser navigation sensor to integrate a BT 2.1 transceiver, stand-alone baseband processor and VCSEL laser surface illuminator into a single package resulting in longer battery life – more than six months using two AA batteries

The BT2.1 protocol ensures fast and secure connectivity when used on Bluetooth-enabled laptops, netbooks and desktops. Avago LaserStream™ technology is best known for its excellent tracking on virtually all surfaces—a very important marketing feature that is attractive to both mobile users and gaming enthusiasts. As shown in Figure 1, the ADNS 7630 combines the latest Bluetooth technology with Avago's optical navigation technology for fast, flexible and convenient mouse design and manufacturing.
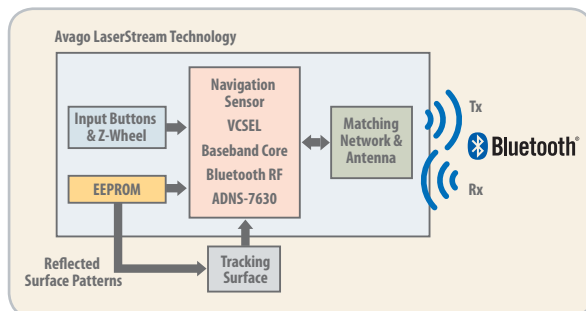

Figure 1: ADNS-7630 implementation of a Bluetooth 2.1 mouse

Value-added Integration – Agility, Flexibility and Easy Design

The ADNS-7630 contains an Image Acquisition System (IAS), a Digital Signal Processor (DSP), and Bluetooth HID stream output. Images acquired by the IAS are processed by the DSP to determine the direction and distance of motion. The DSP generates the Δx and Δy relative displacement values which are then converted to Bluetooth HID data for wireless transmission to the Bluetooth host. All these processes are automatically executed by the ADNS-7630 without the mouse manufacturer being involved with source code or firmware.

Avago's SoC BT 2.1 navigation engine is the first mouse sensor that has a fully synchronized BT 2.1 navigation core

and baseband circuitry. Timing of the sensor measurements and transmission by the radio have been optimized in order to minimize interference. In addition, VCSEL laser diode illumination is also included in the same package.

Avago's free configuration software will configure various settings like Bluetooth Device address, numeric Product Identifier, Vendor Identification, I/O buttons, LED indicator, low power management and others. All the settings are stored in an external 128 kbit EEPROM. The software and easy EEPROM programming allows a mouse manufacturer to produce many different mouse models for different users and geographic markets from one basic ADNS-7630 design.

Mouse Performance and Battery Life

The ADNS-7630 SoC combines Avago's BT 2.1 LaserStream sensor and 100 percent synch capability to allow manufacturers to design feature-rich mice with very long battery life. With integration between the optical sensor, baseband and radio, mouse manufacturers no longer need to worry about firmware debugging and technical support from various parties.

For a typical Bluetooth mouse user, besides reliable Bluetooth connectivity, two other important factors are battery life and wake up latency. A good Bluetooth mouse design is able to strike a balance between these two. The ADNS-7630 design is so flexible that it allows the mouse manufacturer to take charge of their power management strategy by selecting how many sniff modes are supported and by defining each sniff mode in terms of slots and duration. The longer each sniff interval, the better the battery life, but at the expense of wake up time. By enabling sleep mode (similar to disconnect), battery life can be optimized but a longer wake time should be expected, however, most Bluetooth mouse users will find this acceptable. To fully maximize battery life, a manufacturer can also program the mouse to resume connection only through a mouse button click and not through mouse movement.

Battery life indication is a special feature embedded in the ADNS-7630 to eliminate surprises for the end user. For example, a mouse manufacturer can program the low-battery indicator to blink when the battery level is under 2.2 V and to shutdown the mouse when the battery level is lower than 2.0 V. Since a blinking LED will drain power more rapidly the duty cycle, duration and interval of the LED indicator is programmable too.

In general, a design based on Avago's ADNS-7630 will produce a mouse that operates for more than 6 months when powered from 2 AA batteries.

ADNS-7630 LaserStream™ Navigation Sensor and Eye Safety

ADNS-7630 based mice feature up to 30 inches per second (ips) high-speed motion detection and 8 g acceleration detection. In addition, laser power is pre-calibrated to meet IEC/EN60825-1 Class 1 Eye Safety standards. Laser power calibration is not required at the mouse manufacturer's site, greatly reducing assembly time and design cost. As shown in Figure 2, the ADNS-7630 will also shut off the laser whenever a fault is detected, assuring eye safety during the product life cycle.

Figure 2: ADNS-7630 single fault detection and eye safety feature block diagram

Wide Feature Set and Programmability Expand Marketing Options

Avago's ADNS-7630 design carefully considered the needs of mouse manufacturers and markets by supporting multiple I/Os which can be configured as buttons and LED indicators. Each button can be programmed to have several functions dependent on single-clicks double-click, or long pressed. Button functions include KeyMap(KM) for keyboard shortcut keys, media buttons supporting audio control, and mouse resolution increase/decrease/rotate features. Up to ten input buttons are supported. The ADNS-7630 also supports both mechanical and optical z-wheel, and tilt-wheel functions.

For mouse manufacturers who choose to use accept all the default settings provided by Avago, only a few mouse settings need to be set before shipment to the final end customers. As an example, SDP Service Name, Service Description, Provider Name, VID, PID and Bluetooth Address are typical settings manufacturers program into their mice.

Bluetooth BQE Test and Compatibility Issue

The Bluetooth SIG Qualification Enforcement Program is required for every commercial product using Bluetooth wireless technology and is a prerequisite for using the Bluetooth word mark and logos. Avago Technologies has worked with a Bluetooth Quality Expert (BQE) to ensure that ADNS-7630-enabled mice and our IC designs are qualified and interoperate with other Bluetooth enabled products on

the market. Mice designed with the ADNS-7630 SoC sensor were also sent to an Independent Test Lab for Bluetooth Mouse testing and compatibility testing with a variety of peripherals and/or systems.

The ADNS-7630 has passed all the necessary tests, even under extreme voltage ranges (2.7 V and 3.3 V) and extreme temperatures, so as to assure interoperability with other Bluetooth compliant products and systems.

ADNK-7633 Reference Mouse Design Kit

Avago offers a reference design kit, Figure 3, to qualified mouse designers. It includes samples, lenses, documentation, schematic, Gerber files, IGES for 3D assembly, a base plate, LaserStream mouse, EEPROM configuration programming tool and the A7630 production software tool.

Figure 3: ADNS-7633 mouse design kit

Summary

The ADNS-7630 is the first BT 2.1 laser mouse SoC to be introduced. Key features are:

- High speed tracking on virtually all surfaces
- One package solution
  - Bluetooth 1.1/2.1 baseband DSP and transceiver
  - VCSEL LaserStream technology
- High speed motion detection: 30 ips and 8 g
- Programmable resolution: 200-3000 cpi in 250 cpi increments
- Single 3 V supply
  - Two AA cells
  - Long battery life: Over 6 months
- VCSEL laser illumination
- Configurable low power modes
- No laser calibration required
  - Class 1 eye safety
- Up to 10 button inputs with 3D flip programmable to any button
- Vertical scrolling: Mechanical and optical Z-wheel
- 4-axis sensor rotation: 0°, 90°, 180° or 270°

Additional information on the ADNS-7630 can be found at www.avagotech.com.

Contact us for your design needs at: **www.avagoresponsecenter.com/401**

Avago
T E C H N O L O G I E S

Figure 5 A few additional test points can significantly increase test coverage.

er major issues in that their inputs are unobservable and their outputs are uncontrollable. Large memories and analog- and mixed-signal blocks similarly suffer from inputs that are unobservable and outputs that are uncontrollable. The enable pins of tristates are unobservable. Combinational feedback loops also restrict testability, and test-mode values in capture mode may restrict controllability.

Despite efforts from RTL designers, some parts of the design may still not be observable and controllable and may require the insertion of additional test points. Test-coverage analysis at RTL can help detect where to place additional test points and their resulting impact on test coverage. For example, in one design, adding 12 test points increased the test coverage from less than 94% to more than 98% (**Figure 5**). It is easier to add test points in RTL when you fully comprehend the design's intent than in the later stages of implementation.

In deep-submicron designs—those at the 90-nm and smaller nodes—designers worry about transition faults that can occur at normal clock speed. Stuck-at-fault testing, which typically uses a slow test clock, does not detect transition faults. Designers must perform at-speed testing in which system clocks multiplex the test clocks. This step adds a level of complexity for timing closure. At-speed testing also introduces functional closure challenges, such as those that occur when asynchronous clock domains share the same test clock, which could affect the at-speed test coverage. It is therefore critical to estimate at-speed test coverage at RTL and fix potential functional and timing-closure issues.

DFT (design for test) poses a unique challenge for IP reuse. IP that meets test-coverage goals in a previous design could fall short in the current design. For example, if some inputs of the IP are tied to constants in the current design, parts of the IP may become uncontrollable. This issue could affect the test coverage of the SOC. Hence, you must perform test-coverage analysis at both the block/IP level and the SOC level.

DESIGN CONSTRAINTS

Design constraints are a critical part of the design intent. They capture the designer's requirements on the performance, power, and area from implementation. The quality of constraints is just as critical as the quality of RTL in synthesis. At this early stage, designers usually manually define constraints for clock frequencies, input and output delays, modes of operation, and exceptions—false and multicycle paths.

Because this step is the starting point for implementation, the completeness and correctness of constraints are critical in meeting design closure.

You might catch some constraint issues during synthesis if you carefully examine the synthesis transcript. You might find, for example, missing-clock or -mode constraints when the design is using a multiplexed clock (**Figure 6**). Other constraints may occur when input and output delays reference an incorrect clock, multiple mode constraints tie the same node to conflicting constant values, or timing exceptions are missed on asynchronous clock-domain crossings.

However, synthesis or static-timing analysis may not catch more serious issues in design constraints. These issues typically involve constraints that are simply assertions to the synthesis and static-timing analysis; they thus remain undetected. You may not catch such issues until final chip integration or, worse yet, in silicon.

For example, a generated clock does not derive from the declared source clock, and the waveform for a generated clock differs from the implied waveform, depending on the presence and location of inverters in clock-divider circuits. Other examples include missing clock latency or uncertainty, missing delay constraints on primary inputs or outputs, block-level constraints that are more relaxed than the chip-level requirements, incorrect or insufficient timing budget along a snaking path, or an incorrect multiplier for multicycle paths.

> + **Go to the Talkback section at www.edn.com/ ms4369 to post a comment on this article.**

Analysis of design constraints in RTL design can help avoid these issues. Designers can exploit this analysis to generate the clock, input, and output constraints in a correct-by-construction way, thereby eliminating many overlooked bugs. For instance, you can use clock-domain-crossing analysis and knowledge of asynchronous control signals to generate timing exceptions. At RTL, you can address inconsistencies between block- and chip-level constraints by comparing the sets of constraints in the context of the complete RTL design.
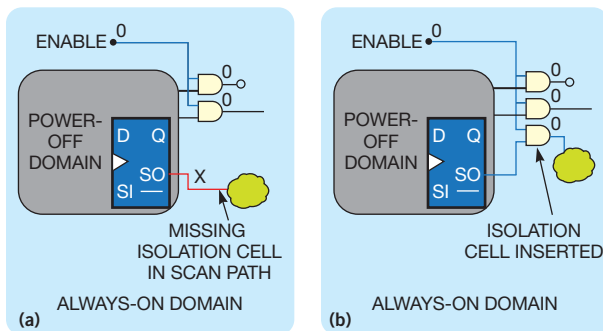
POSTSCAN-NETLIST QUALITY

At the postsynthesis stage, the design has gone through logic synthesis, resource sharing, Boolean optimization, and scan-chain insertion. Assuming that RTL and constraints in

Figure 6 This situation requires a mode constraint to propagate the clock.

Figure 7 A power-domain bug in the scan path occurs (a), and you can repair it using another approach (b).

synthesis are high-quality, the resulting netlist should be in good shape. However, a lot has changed in the design, and constraints may also have changed. You can take some measures to ensure quality at the postsynthesis, postscan netlist and constraints.

At this stage, the design may contain power and voltage domains to manage power consumption. You now must perform an exhaustive verification of these domains to ensure proper insertion of level shifters and isolation logic. Even if you inserted these level shifters and isolation logic before synthesis or scan-chain insertion, the design may still need updates. In one such example, an isolation cell is missing in the scan path because DFT designers inserted the scan logic after the implementation of power domains (**Figure 7**). Adding the isolation cell in the scan path fixes a potential functional bug or a high-leakage path to ground.

Such power-domain bugs also commonly occur when designers forget the required "don't-touch" constraints in the synthesis script. This omission can cause buffer optimization during synthesis, which removes level shifters or isolation-logic cells.

SOC designs are now so large and complex that implementation may need to be hierarchical. Designers typically perform synthesis at the block level, and chip integration occurs at the netlist level. In such scenarios, you must merge block-level-synthesis constraints into chip-level constraints. Manually merging constraints can be error-prone and could lead to an incorrect set of chip-level constraints for later timing closure. You could at this stage apply constraint consistency and correctness checks. You can automatically merge constraints by using block-level constraints and the full-chip netlist to prevent bugs.

If constraints have changed at this stage, it is important to establish equivalence with the original constraints for synthesis. Just as you can perform logic-equivalence checking for design stages from RTL to netlist, you can now establish equivalence between RTL and netlist constraints. Designers can ensure the integrity of design constraints and eventual design closure by adopting constraints equivalence in their flow.

If the implementation flow is hierarchical, you might want to stitch together chip-level test logic at the netlist level. In such scenarios, ensure that the global test clocks and test-

mode signals propagate to individual blocks. You should perform quality checks to ensure that required values propagate to pins on subblocks when they specify a driving condition at a primary input or an internal node of the chip. Similarly, designers can benefit from connectivity checks to ensure that a path exists between two user-specified nodes in the design with an optional sensitization condition (**Figure 8**). For example, imagine that Pin A in Block 1 does not connect to a primary output and that Pin B in Block 2 is observable. By establishing a path between Pin A and Pin B, you can now ensure that Pin A is observable.

POST-TIMING-NETLIST QUALITY

At the post-timing stage, you must ensure that the design meets timing requirements and starts analyzing timing violations from static-timing analysis. This stage is another critical one. Timing closure can be a challenge if the design is overconstrained or incorrectly constrained. Other causes of

> **POWER-DOMAIN BUGS COMMONLY OCCUR WHEN DESIGNERS FORGET THE REQUIRED "DON'T-TOUCH" CONSTRAINTS IN THE SYNTHESIS SCRIPT. THIS OMISSION CAN CAUSE BUFFER OPTIMIZATION DURING SYNTHESIS.**

problems could be structural defects, such as combinational loops, excessive levels of logic, or unregistered outputs from blocks or IP, all of which you should have detected at earlier stages of design.

Timing exceptions fall into two broad categories: false paths and multicycle paths. False paths between two registers are those you cannot sensitize in the design or are otherwise irrelevant for timing closure. Multicycle paths, on the other hand, are possible but take multiple clock cycles to complete. Unless you identify the false and multicycle paths in the design constraints, static-timing-analysis tools assume that all paths are possible and single-cycle.
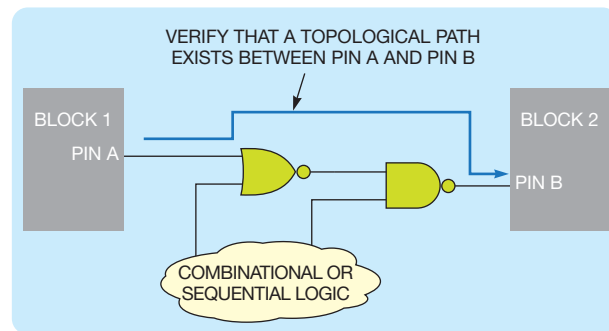
Figure 8 When you assemble IP at the netlist level, it's best to do connectivity checks between user-defined nodes.

<table>

Design	No. of false paths	Worst negative slack (nsec)	Total negative slack (nsec)	No. of instances
Block 1	None	−5.462	−16.37	118,426
With additional false paths	595	−4.073	−13.829	119,162
Impact		−24.4%	−15.5%	0.6%
Block 2	None	−6.019	−11.522	77,064
With additonal false paths	1242	−2.227	−6.39	77,166
Impact		−63%	−44.5%	0.1%

TABLE 1 IMPACTS ON TIMING CLOSURE

An incorrect timing exception can lead to a critical timing failure in silicon. On the other hand, every timing exception that remains unidentified is unnecessary and represents a wasteful timing-closure burden. It is therefore a fine balancing act to find just the right timing exceptions. You must at least formally verify all timing exceptions in use to ensure their validity.

Another step that could accelerate timing closure is to look for additional timing exceptions, especially in those paths that are violating timing. You should formally verify every such path as a possible false- or multicycle-path candidate; if it is false or multicycle, you should add it to the list of timing exceptions for static-timing analysis. Consider the results of timing analysis on two timing-critical blocks from a multimedia design (**Table 1**). The timing results dramatically improve when you identify additional timing exceptions from paths that initially failed to meet timing. The impact on gate count and, hence, area is minimal.

POSTPLACEMENT-NETLIST ANALYSIS

By the postplacement-analysis stage, the design has entered physical implementation and has undergone physical synthesis, placement, and clock-tree synthesis. You should repeat quality measures on the now fully placed netlist. You now have a more accurate assessment of the power, area, timing, and test coverage, and you can compare this estimate with the estimated results from RTL to identify blocks that may have deviated.

You can at this stage perform additional netlist-quality checks, such as floating pins or nets; clock, select, enable, or reset pins that tie to constants; unused or disabled cells; undriven or multiply-driven nets in the netlist; overloaded cells; underloaded cells, wasting area and power; pins connected to specific nets, such as tristate, clocks, and resets; scan-chain nets with more than the maximum number of elements; and high-leakage or snake paths. You should also check that pins connecting to the same net are of the same connectivity class.

Before clock-tree synthesis and before a clock network exists, you should specify, in the design constraints, the values the tools should assume for clock latency and clock slew rate. However, assuming that you've inserted the clock tree by this stage of the design, it is time to compute and apply the delays and slew. At this time, you must also update and verify design constraints for two critical areas. First, replace assumed transition times on individual flip-flops with transition times only on the primary inputs to the design. Second, set clock delays to propagated, rather than to a user-defined network latency on clocks.

POSTROUTE-NETLIST ANALYSIS

Postroute-netlist analysis represents the home stretch for design implementation, yet design teams spend a lot of time and effort at this stage to close timing, signal integrity, manufacturability, power integrity, and a host of physical effects. Assuming that you've followed the quality measures in earlier stages, the design and constraints should be fairly high quality, and you should focus on these physical effects. In addition, you should place significant effort in layout and physical verification, tackling process variability and other manufacturing issues. This stage also encompasses final sign-off on power, timing, testability, and die size; hence, it is best to repeat the quality measures from earlier steps as part of the final sign-off.

In short, the quality of a design and its associated constraints have a large impact on design closure. You can, however, take a series of quality measures to improve the chances of design closure. It is also important that you take most of these measures at the early stages of design, especially at RTL, at which point you best comprehend the user's intent. The later in the implementation flow that you address design quality, the less impact it is likely to have on design closure. If you get design goals and quality objectives right from the start, it is only a matter of staying the course during implementation.EDN

AUTHORS' BIOGRAPHIES

Piyush Sancheti is senior director of business development at Atrenta Inc, where he is responsible for Atrenta's strategic alliances with key members of the semiconductor supply chain. Sancheti has more than 15 years of experience in various marketing and product-management roles in the EDA industry at Sequence Design, Senté, and Cadence. He has a master's degree in computer engineering from Iowa State University (Ames, IA). You can reach him at psancheti@atrenta.com.

Sanjay Churiwala is senior director of engineering at Atrenta Inc, where he has worked for eight years. His responsibilities include EDA-tool development and application responsibility for strategic customers. He has a bachelor's degree in technology from the Indian Institute of Technology (Mumbai). You can reach him at sanjayc@atrenta.com.

Rob Knoth is senior product manager at Magma Design Automation, where he has worked for three years. He has a bachelor's degree in electrical engineering from Purdue University (West Lafayette, IN). You can reach him at rknoth@magma-da.com.

Low EMI DC/DC µModule Family

VERIFIED TO POWER
Multigigabit Transceivers

V_IN
4.5V to 28V

LTM4606

V_OUT @ 6A
0.6V to 5.5V

Qualified for Fast I/O SERDES, Transceivers

Meets CISPR 22 Class B (Radiated Noise)

CISPR22, CLASS B

Emissions Level (dBµV/m)

Frequency (MHz)

Meets CISPR 25 Level 5 (Conducted Noise)

CIS25QP

Signal Amplitude (dBµV)

Frequency (MHz)

Simple Power Conversion for Noise-Sensitive Systems

Linear's new family of DC/DC µModule™ regulators is designed for noise-sensitive systems that require lower EMI to meet radiation standards or lower BER (Bit Error Rate) in multigigabit/second data transmission systems. Noise is attenuated with onboard input filters, controlled slew rate, shielded inductors and careful layout. This system-in-a-package regulator family has been certified for radiated and conducted noise. The LTM®4606 has passed jitter tolerance tests for high speed data transmission in SERDES applications.

▼ Low Noise DC/DC µModule Regulators

Part Number	V_IN Range	V_OUT Range	I_OUT	EMI	SERDES, High Speed I/O	LGA Package
Narrow Voltage Range						Pin Compatible 15mm x 15mm x 2.8mm
LTM®4606	4.5V to 28V	0.6V to 5.5V	6A	CISPR 22 Class B	Powers Multigigabit Transceivers	
Wide Voltage Range						
LTM4612	4.5V to 36V	3.3V to 15V	5A	CISPR 22 Class B	N/A*	

*I/O voltages require less than 2.5V. See LTM4606. Use LTM4612 for low noise intermediate bus voltages (≥ 3.3V).

LINEAR TECHNOLOGY

designideas

READERS SOLVE DESIGN PROBLEMS

Eight-digit counter works with common anode or common cathode

Charaf Laissoub, Valeo Interior Controls, Créteil, France

A classical design for directly driving eight seven-segment LED displays requires as many as 15 I/O lines. Previous Design Ideas have described many approaches for using a maximum number of LEDs with a minimum number of I/O lines (**references 1** through **5**). The following idea reuses one of these approaches to drive a maximum number of seven-segment LED displays, and it may be useful in designing a low-component-count, low-power, and low-cost LED-display module for a 24-bit frequency meter, for example.

You can use the circuit in **Figure 1** to replace classical designs for digital counters that use TTL (transistor-transistor logic) or CMOS ICs. The single microcontroller is less expensive and readily

available. By using conditional assembly in your programming, you can choose between common-anode and common-cathode configurations.

The algorithm uses double multiplexing, driving one digit at a time, segment by segment. This technique suits battery-powered designs because the circuit consumes a constant current of less than 2 mA when using superbright, seven-segment LED displays, such as KingBright's (www.kingbright.com) SC52-11EWA, and 270Ω resistors R_0 to R_7. Assembling the eight digits, DS_7, DS_6, DS_5, DS_4, DS_3, DS_2, DS_1, and DS_0, on a PCB (printed-circuit board) involves linking their corresponding pins A7, B6, C5, D4, E3, F2, G1, and CA0 to the I/O line, R_{B0}. **Figure 1** shows the connections.

This circuit uses the Microchip (www.microchip.com) PIC16F628A for test purposes. You can download assembly code at www.edn.com/100715dia and use it with any PICmicro midrange family, providing that a full 8-bit port is available.EDN

NOTE: CHOOSE EIGHT COMMON-ANODE OR EIGHT COMMON-CATHODE, SEVEN-SEGMENT LED DISPLAYS AND UNCOMMENT ONE OF THESE TWO LINES IN THE INCLUDE FILE: "8dgtcacc.inc" #define use CAdisplay IF USING COMMON-ANODE, SEVEN-SEGMENT, DISPLAY OR #define use CCdisplay IF USING COMMON-CATHODE,SEVEN-SEGMENT DISPLAY.

Figure 1 This circuit can replace classical designs for digital counters that use TTL or CMOS ICs.

designideas

REFERENCES

[1] Anonymous, "Microcontroller provides low-cost analog-to-digital conversion, drives seven-segment displays," *EDN*, May 10, 2007, pg 80, http://bit.ly/b3KBre.

[2] Raynus, Abel, "Squeeze extra outputs from a pin-limited microcontroller," *EDN*, Aug 4, 2005, pg 96, http://bit.ly/aKKGEf.

[3] Jayapal, R, "Microcontroller's single I/O-port line drives a bar-graph display," *EDN*, July 6, 2006, pg 90, http://bit.ly/barbMI.

[4] Lekic, Nedjeljko, and Zoran Mijanovic, "Three microcontroller ports drive 12 LEDs," *EDN*, Dec 15, 2006, pg 67, http://bit.ly/deQTib.

[5] Gadre, Dhananjay V, and Anurag Chugh, "Microcontroller drives logarithmic/linear dot/bar 20-LED display, *EDN*, Jan 18, 2007, pg 83, http://bit.ly/azQ9dw.

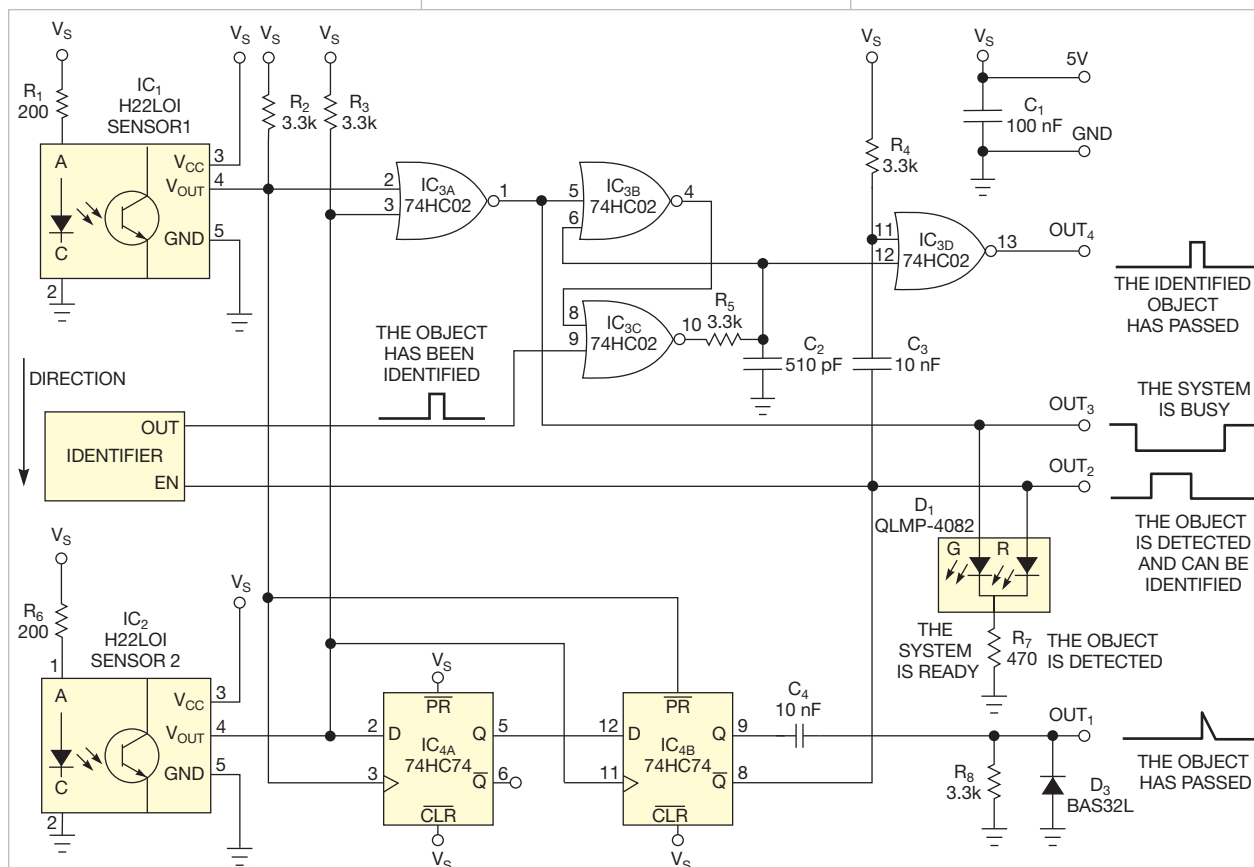# Count objects as they pass by

Vladimir Rentyuk, Zaporozhye, Ukraine

Counting objects is easy when the objects have regular motion, but it becomes more difficult when the objects vibrate or you have to manually move them. To deal with this problem, you need a reliable, error-free system of detecting an object. In this case, simple circuits with optical interruption switches, IR barriers, or other sensors don't work because objects may cross a sensor more than once.

The circuit in **Figure 1** solves this problem for objects such as cards, boxes, and even people. You can detect other objects if you use the proper sensors. The circuit produces two control pulses. One, OUT_1, occurs if any object runs through this system. It produces another pulse, OUT_4, only if an identified object passes through the system. These pulses let you count both the number of objects and the identified objects that pass through the system. The system doesn't produce counting errors if objects repeatedly cross

THE CIRCUIT PRODUCES TWO CONTROL PULSES.

each sensor, even when the objects return to the system, provided the object does not move from Sensor 1's zone. The system requires no difficult mechanical unit for stabilizing the speed of the moving objects. The counter selects objects that are slightly longer than the distance between its sensors. **Figure 2** shows how the circuit tracks an object between the sensors.

Figure 1 This circuit counts objects such as cards, boxes, and even people.

Feel the boost

Boosted amplifier blasts twice the watts of a conventional amplifier

The MAX9730 is a mono speaker amplifier with an integrated charge pump that delivers 2.4W into 8Ω. This integration eliminates the need for a boost converter and large inductor, thus minimizing total solution size.

Increased output power with small footprint

- **2.4W output power from a standard 5V supply (1W from 3.3V supply)**
- **Class G architecture improves efficiency and eliminates EMI concerns**
- **Ultra-small, 2mm x 2.5mm solution**

Part	Supply (V)	P_{OUT} into 8Ω at 1% THD+N (W)	Output Voltage Swing (V_{P-P})	PSRR (dB)	THD+N (%)	Boost Type	Speaker Type
MAX9730		2.4	—	77	0.007	Capacitor	Dynamic
MAX9738	2.7 to 5.5	—	16	60	0.004	Inductor	Piezoelectric
MAX9788		—	20	77	0.002	Capacitor	Piezoelectric

www.maxim-ic.com/MAX9730-info

For free samples or technical support, visit our website.

Figure 2 The circuit tracks movement using an object's length and direction.

TABLE 1 SIGNAL DESCRIPTION	
Trace or zone	Indication
Trace 1	OUT$_1$: common counting of passing objects
Trace 2	OUT$_2$: detection of an object
Trace 3	OUT$_3$: ready to take an object
Trace 4	OUT$_4$: counting an identified object
Zone A	Object identified
Zone B	Object not identified

The object counter comprises IC$_1$ and IC$_2$, H22LOI sensors from Fairchild Optoelectronics Group (www.fairchildsemi. com). The sensors are optical-interruption switches with open-collector outputs. A control circuit uses IC$_4$, a 74HC74 dual D-type flip-flop, and IC$_3$, a quad, two-input NOR gate. Sensors IC$_1$ and IC$_2$ produce logic-low output levels when objects are not between the IR diode and the IR receiver. LED D$_1$ shines green if the system is ready to take an object and red if the system has detected the object.

If the system is busy, LED D$_1$ does not illuminate. Resistor R$_5$ and capacitor C$_2$ protect RS trigger IC$_{3B}$/IC$_{3C}$ from the chance of failure due to undesired signals. Some sensors require that you add buffers between them and the control circuit.

Table 1 describes the circuit's signals. A downloadable sheet describing the sequence of events is available at www. edn.com/100715dib.EDN

Modified DDS functions as baud-rate generator

Vardan Antonyan, Glendale, CA

You can often use an available oscillator to generate a baud-rate clock for a UART. You must divide the oscillator frequency to attain the proper baud rate, but dividing can produce baud-rate errors. **Table 1** shows the percentage of error when you generate a baud rate using an 8-MHz crystal oscillator and a conventional binary divider. The system in this Design Idea obtains a clock 16 times faster than the baud rate.

Errors in baud-rate setting increase when the oscillator frequency doesn't match. In this case, you can add an oscillator operating at 18.432 MHz, for example, to minimize the error rate. Alternatively, you can use DDS (direct digital synthesis) to reduce errors at higher baud rates using the same oscillator (**Table 2**).

Reference 1 describes basic DDS operation. This design uses a simpler version of DDS with only a square-wave output (**Figure 1**). You can extract the square-wave output from the MSB of the phase accumulator. You can also add the divide-by-two

TABLE 1 BAUD RATE WITH REGULAR DIVIDER		
Baud rate	Divisor	Error (%)
50	10,000	0
300	1666	0.04
600	833	0.04
2400	208	0.16
4800	104	0.16
9600	52	0.16
19,200	26	0.16
38,400	13	0.16
57,600	8	7.84
115,200	4	7.84
230,400	2	7.84

stage to make the resulting signal with a 50% duty cycle. Calculate the baud-rate clock frequency using baud-rate clock= (reference clock×tuning word/2^N)/2, where N is the number of bits for the phase accumulator. A Verilog implementation of the DDS baud-rate generator using a 20-bit phase accumulator and 16-bit tuning word is available at www. edn.com/100715dic.EDN

REFERENCE

[1] "A Technical Tutorial on Direct Digital Synthesis," Analog Devices, 1999, www.analog.com/static/imported-files/tutorials/450968421DDS_Tutorial_rev12-2-99.pdf.

TABLE 2 BAUD RATE WITH 20 DDS BITS		
Output frequency	Phase word	Error (%)
50	13	−0.825
300	78	−0.825
600	157	−0.182
2400	629	−0.023
4800	1258	−0.023
9600	2516	−0.023
19,200	5033	−0.003
38,400	10,066	−0.003
57,600	15,099	−0.003
115,200	30,198	−0.003
230,400	60,397	−0.002

Figure 1 You can generate a baud-rate clock from an available oscillator.

High-voltage integrated pulsers and switches save space and reduce power consumption

Ideal for portable ultrasound equipment applications

TRANSDUCERS

- No need for external regulators
- Reduces system power consumption
- Smallest packages in the industry

Improvements over the competition

Quad-channel, high-voltage, 2.0A digital pulsers (MAX4940/MAX4940A)

- Integrated active clamps minimize 2nd harmonic distortion and reduce power consumption by 66% vs. HDL6V5581
- Capacitive architecture eliminates need for large, costly floating regulators and complex power sequencing
- AC-coupled architecture reduces overall footprint by up to 45% vs. HDL6V5581
- Five-level pulsing with active damping (MAX4940A) enables transmission for complex systems

16-channel, high-voltage switches (MAX14800*/MAX14801*/MAX14802/ MAX14803)

- Ultra-small, 7mm x 7mm CSBGA package saves board space
- Integrated overvoltage protection diodes prevent damage caused by high-voltage overshoots (MAX14802/MAX14803)
- Designed with larger, high-voltage supply range for better linearity
- 33% faster turn-on/off times provide fast frame rates for improved imaging

*Future product—contact factory for availability.

www.maxim-ic.com/Ultrasound-info

For free samples or technical support, visit our website.

DC-voltage doubler reaches 96% power efficiency

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

The voltage-doubler circuit in **Figure 1** can convert 2.5V dc to 5V dc or 1.8V to 3.3V. Most voltage doublers use an inductor, but this circuit doesn't need one. The circuit uses a capacitor, C, by charging it through serially connected switches. The charge switches let capacitor C charge, and the discharge switches are open. In the subsequent discharging phase, the charge switches are off, and the discharge switches close. The two discharge switches now connect capacitor C between the source of the input voltage, V_S, and the output capacitor, C_{OUT}. This connection scheme lets the applied voltages combine. Thus, the volt-
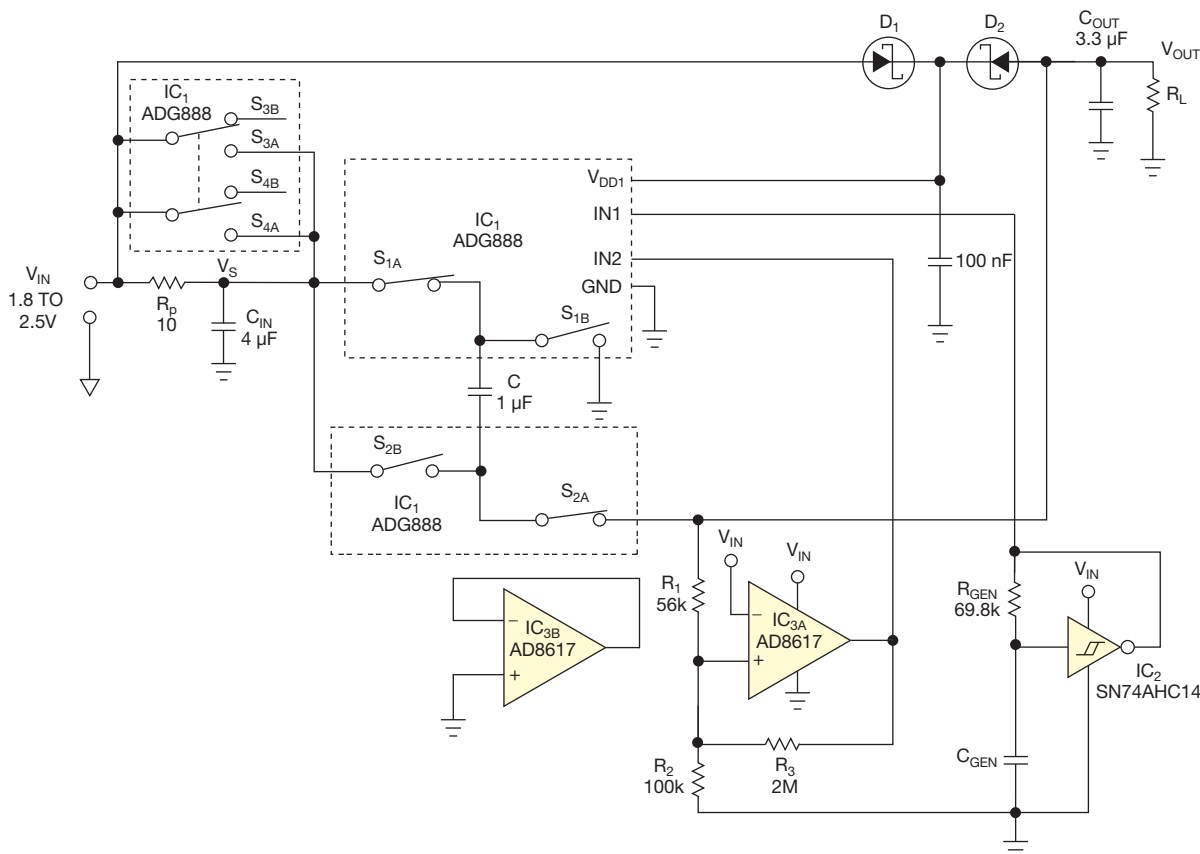
age at the output terminal has a value close to $2V_S$.

The two phases of operation repeat periodically at frequency f, which clock generator IC_2 determines. The duty cycle is about 50%, but the value isn't all that critical. One half of the Analog Devices (www.analog.com) high-performance ADG888 analog multi-switch provides the switching. The IC's two halves have independent control, so the other half occasionally shorts R_P, the 10Ω inrush-current-limiting resistor, which protects the charge switches from an initial overcurrent. That current occurs after power-on, before the output voltage reaches the predetermined per-

centage of the output's full voltage.

A micropower op amp, IC_{3A}, runs as a comparator with hysteresis. It compares input voltage to output voltage. Its output starts low and then goes high, which turns on paralleled switches S_3 and S_4. The comparator's action is ratiometric because the reference input voltage at the inverting input is the input-supply voltage, V_{IN}. This connection is possible because of the AD8617's rail-to-rail input/output operation. The circuit also provides overload protection for an excessive load, which connects to the circuit's output before power-on.

During soft start, the output voltage can't reach the threshold level for loads below a certain value. Consequently, the circuit remains in soft-start mode. The minimum value of R_L, which activates the protective subcircuit, is $R_L \leq m^2 \times (\alpha/(1-\alpha)) \times R_P$, where the multiplication factor $m = (V_{OUT}/V_{IN})$ and α is

NOTE: C_{GEN} HAS AN OPTIMAL VALUE OF 62 pF WHEN R_L IS APPROXIMATELY 180Ω.

Figure 1 You can use this step-up dc/dc converter in applications in which power efficiency is a critical issue.

a fraction of V_{OUT}, at which the soft start turns off. For m=2, α=0.8, and R_P=10Ω, R_L is 160Ω. Thus, loads of 160Ω or less will overload the circuit if you connect them to the circuit's output before power-on. IC_2 and IC_3 get their power from the input supply. IC_1, however, switches voltages of as much as $2V_{IN}$, and its V_{DD1} supply-voltage pin must remain at the same level. An analog OR switch comprising Schottky barrier di-

odes D_1 and D_2 provides that voltage. The higher of the input or output voltages appears at the V_{DD1} pin of IC_1. The high levels of output voltages for both IC_2 and IC_3 suffice for control of IC_1 because the ADG888's data sheet allows a $0.36V_{DD1}$ value for the high value at the control inputs. The circuit has been tested at an input voltage of 2.386V, R_L of 178.46Ω, a frequency of 200 kHz, a supply voltage of 2.377V, an input sup-

ply current of 51.285 mA, and an output voltage of 4.588V. Evaluating these data gives a multiplication factor of 1.929 and power efficiency of 96.39%.

This power efficiency remains more than 96% for frequencies of 150 to 350 kHz. The 9-mV drop at the switch-shorted R_P at the given input current indicates that the on-resistance of the paralleled switches has a value of approximately 0.175Ω.EDN

Microcontroller's serial port measures pulse width

Vishwas Vaidya, Tata Motors Ltd, Pune, India

Many industrial and instrumentation systems need to measure the duration of pulse inputs, such as frequency from rotational-speed sensors, gating and strobe pulses from external systems, and PWM (pulse-width-modulated) inputs. Designers generally use on-chip timers and edge-driven interrupts for this purpose. If one of these components is unavailable, however, you can employ an unused on-chip serial-synchronous receiver to make those measurements.

You can set the baud rate of the serial-port receiver for the necessary timing accuracy. The receiver interrupts the microcontroller after every 8 bits. You can embed the pulse-width acquisition routine, which resides in your application program, to read the byte that the ISR (interrupt-service routine) receives. It counts and accumulates the number of ones and zeros the bytes receive to measure the duration of an incoming pulse (**Figure 1**).

The algorithm measures the duration between two consecutive rising edges.

The microcontroller detects a rising edge or a falling edge when a received byte is neither 0xff nor 0x00. If a byte is less than 0x80H (100 000b), then the byte marks a rising edge. If the byte is equal to or greater than this value, then the byte marks a falling edge.

The leading rising edge resets the bit counter to the number of trailing ones in the rising-edge byte by shifting the bits to the right. The bit counter increments by eight at the arrival of every byte, including the one that marks the falling edge. When the counter receives the trailing rising edge, marked by the next rising-edge byte, which is greater than 0x80h but less than 00H, it again counts the number of leading zeros in this byte and adds them to the accumulated-bit coun-

ter. The accumulated-bit count at this point directly relates to the time period of the pulse train by a factor equaling the baud rate.

Figure 2 depicts a 200-Hz pulse train, which has a 5-msec period between two consecutive rising edges. The baud rate is 256 kbaud. During a measurement cycle, assume that the leading rising edge is marked as 0011 1111b. The microcontroller counts the number of trailing ones by shifting them right and initializing the bit counter as six. This count corresponds to approximately 23.43 μsec.

Next, every byte before the rising-edge byte increments the bit counter by eight. Simple calculation shows that the sum is 159 bytes, or 1272 bits. At this point, the total bit count is 1278, including six one bits received in the first rising-edge byte.

The pulse train now encounters its trailing rising-edge byte as 0011 1111b. When this encounter occurs, you need to shift the zeros left to count two bits. The total bit count between the rising edge now is 1280. At a 256-kbaud rate, this figure corresponds exactly to 5 msec, or 200 Hz.

Figure 3, a flow chart, is available with the online version of this Design Idea at www.edn.com/100715did. It explains how you can use this concept to measure frequencies in hundreds of hertz.

You can tailor this bit-counting concept to your application's requirements. For measuring only a low period of a pulse, you need to detect a falling edge and count the bits until you encounter a rising edge. You can use this concept to read an incoming PWM signal by reading high periods of a known incoming pulse frequency.EDN

Figure 1 A microcontroller's serial port can receive pulses from various sources.

TOTAL 1280 BITS AT 256 kBAUD

159 BYTES (1272 BITS BETWEEN RISING EDGES)

BIT PATTERN 00111111b FOR LEADING RISING EDGE

BIT PATTERN 00111111b FOR TRAILING RISING EDGE

Figure 2 The algorithm uses a 001111 bit pattern to detect edges.

supplychain
LINKING DESIGN AND RESOURCES

Outsourced engineering services show solid growth

Last year's economic environment, in which customers needed to continue driving their key products to market with fewer resources, spurred an increase in outsourced engineering, says Andrew Femrite (**photo**), manager of Arrow Electronics Inc's (www.arrow.com) Engineering Solutions Center. "Hiring has not yet increased dramatically, so outsourcing continues not only as a short-term tactic but also as a mainstay," he adds.

The economic downturn also drove work to independent design companies. "Many companies downsized over the last few years, but outsourced-engineering companies did well," says Adam Tavin, chief executive officer of Focus Product Design Inc (www.focuspdm.com), a design-engineering company. "The work still needed to get done, even though companies had downsized."

Analysts note that companies in the electronics industry have recently become comfortable with outsourcing. "In the past three years, more work is getting outsourced," says Mukesh Dialani, research manager at IDC (www.idc.com). "At some level, confidence has grown in what engineering vendors have to offer. Outsourced engineering has been validated with successful outsourcing engagements."

Outsourced design is attractive in a number of areas. Some companies use outsourcing to keep old products alive while their in-house team focuses on new products. "We saw clients who kept their legacy products going for a small investment in outsourced engineering," says Rafael Cruz, vice president of engineering services at Avnet Inc (www.avnet.com). "The new engineering let them continue to derive revenue from legacy products instead of obsoleting them."

For some companies, including large manufacturers in the United States, design companies have become sources of specialized expertise. Companies are also turning to outside design for expertise they need but don't have sufficient demand to build in-house. "Companies are wrapping new technologies, such as lighting or wireless capability, around their core competence, and they see outsourcing as a means for quickly adding a new competence from an expert," says Arrow's Femrite. "Our customers are focusing much more on the competency of their design partners than having them in their own yard."

Companies are outsourcing anything from entire products to pieces of designs, such as boards or software. "We see outsourcing needs for a broad range of items, from proof of concept to compliance testing to software development to system-on-module customization to ASIC design," says Femrite. "Most of the time, companies pursue third-party-design support for pieces of their system outside their core competence. We see fewer full-turnkey design contracts."

Emerging technology is also a candidate for outsourcing. For many companies, it's more cost-effective to buy cutting-edge technology rather than hire for it. "We're seeing a lot of growth in 'green' engineering," says John Myung, vice president of sales and marketing at Focus Product Design. "In automotive, companies are looking for new battery design. We're also seeing more design work for companies analyzing solar panels."

Companies doing outsourced-design engineering range from distributors and component suppliers that support their customers' product development to specialty design services. "Specialty-design houses [are] specific to different sectors, [such as] medical or industrial," says Eric Miscoll, a principal at Charlie Barnhart and Associates (http://charlie barnhart.com). "What they have in common is their expertise in building products."

—by Rob Spiegel

OUTLOOK

GARTNER EXPECTS RECORD SEMI SALES THIS YEAR

Analysts at Gartner Inc (www.gartner.com) have forecast that 2010 will show 27.1% revenue growth on 2009's $228 billion for $290 billion in revenue. The company expects 2011 revenue to reach $307 billion and for growth to continue through 2014.

"Sequential semiconductor growth has been strong over the last five quarters—well above seasonal norms, and manufacturing capacity is tight," says Bryan Lewis, research vice president at Gartner. "Chip-revenue growth is clearly outpacing system-revenue growth."

The PC and mobile-phone markets will account for about 40% of the semiconductor market's growth in 2010. Processors' average selling prices are firming up, and 2010 PC processor revenue should grow 15.5%, Gartner estimates. Gartner in May forecast that PC shipments will climb 22% this year.

"Strong PC growth and rising DRAM prices are causing the 2010 DRAM market to surge 78%, making it the strongest-performing semiconductor-device market," Gartner reports. "Analysts believe that the demand for media tablets, such as Apple's iPad, will noticeably impact the PC market by 2013, further fueling growth."

—by Suzanne Deffree

productroundup

OPTOELECTRONICS AND DISPLAYS

LED driver ups HB-LED reliability

The ZXLD1370 LED-driver controller increases the reliability of HB (high-brightness) lamps in automotive-, industrial-, and commercial-lighting systems. Operating in buck, boost, and buck-boost modes, it combines control loops and high-side current sensing to ensure accurate current control of LED strings. With a 6 to 60V operating input-voltage range and a typical 1% output-current tolerance, the driver can deliver the high current levels and tight interlamp luminance match that HB-LED systems require. A dedicated external thermistor input and fault-diagnosis outputs, which report the status of the LED driver and the load, achieve active LED thermal management, improving LED reliability and longevity. The ZXLD1370 uses an external MOSFET and has 1-MHz switching frequency and a 100-kHz clamp, reducing the size of the internal inductor and requiring simpler design of any required input EMI filters. The controller sells for $1 (1000).
Diodes Inc, www.diodes.com

LED drivers tout 92% peak efficiency

The charge-pump-based, 180-mA, six-channel FAN5701 and FAN5702 LED drivers backlight TFT LCDs in mobile electronics and feature 92% peak efficiency. The FAN5701 uses two PWM inputs to control the brightness of an array of four and two LED outputs. The configurable FAN5702 has an I²C interface. Both devices' outputs support 64 logarithmic-dimming steps, eliminating distortion, especially at low-brightness levels. They also feature dynamic backlight control. The FAN-5701 and FAN5702 sell for 72 and 95 cents (1000), respectively.
Fairchild Semiconductor, www.fairchildsemi.com

Backlight drivers boost efficiency by 12% in portable devices

The MAX8847 and MAX8848 backlight drivers for as many as six or seven white LEDs come in 3×3-mm TQFN packages. The backlight drivers employ negative-charge-pump architecture to eliminate inline resis-

tance from the battery to the LEDs. As a result, mode switching decreases from one to 1.5 times during battery discharge, minimizing dropout. The devices' adaptive-mode-switching topology individually controls each LED. A serial-pulse/PWM interface controls each LED, enabling multizone light management and allowing as many as 32 brightness levels. A 120-μA quiescent current facilitates always-on displays that require minimal power. Prices for the drivers start at $1.95 (1000).
Maxim Integrated Products, www.maxim-ic.com

Solid-state driver module targets residential, commercial, industrial LED fixtures

The LXMG221W-070034-D0 power-supply module supports five to 16 LEDs and has a universal input-voltage range of 90 to 305V ac, enabling operation in 100, 120, 220 to 240, and 277V-ac, 50- and 60-Hz systems. The Class 2 isolated module provides 90% peak power-conversion effi-

ciency and delivers nonflickering dimming to 10% using 0 to 10V dimming controls and potentiometers. Other features include a power factor of less than 0.9; a total harmonic distortion of less than 20%; a constant-current, single-string output of 700 mA; and a 14 to 48V-dc output for LED loads of as much as 34W. CE and UL1310 certifications are pending, and the module comes in an IP66-rated plastic package. Samples are available for $60.

Microsemi Corp,
www.microsemi.com

Dimmable LED-controller IC suits dimmable and nondimmable LED lamps

↘ Targeting high-power LED lighting, including PAR20, 30, and 38, the SSL2103 ac/dc LED-controller IC finds use in both isolated and nonisolated, dimmable or nondimmable LED lamps. The IC also finds use in 100, 110, and 230V-ac networks and form factors for E27, GU10, or PARx lamps. The SSL2103 sells for $1.35 (1000).

NXP Semiconductors,
www.nxp.com

Touch panel comes in a variety of sizes for wide applicability

↘ The newest addition to the vendor's Crystal Touch capacitive-touch panel line comes in 4.3-, 7-, and 10.1-in. sizes with multitouch capabilities. The Crystal Touch has a layer of glass on the exterior, increasing durability and allowing the panel to handle a wide range of temperatures. Previous touch panels, such as resistive-touch interfaces, use a film on the screen's exterior, which scratches easily and has a shorter life. The panel also provides a clean surface for the secure attachment of cover plates, increasing readability and durability for handling harsh cleaners. The 10.1-in. touch panel sells for $57 (volume quantities).

Ocular LCD Inc, www.ocularlcd.com

Stopped in our tracks

Back in the '70s, I worked for a medical-electronics company that made patient-monitoring equipment. My group developed a minicomputer-based system that accepted bedside data and provided patient status and charting information on a couple of CRTs at a central nurses' station. The equipment comprised a 19-in. rack of interface electronics and a minicomputer. Most of our installations went into hospitals, but it was sometimes difficult to find a place in the hospital for this rather large rack of equipment. Most of the time, we put the equipment into an out-of-the-way closet.

We were feeling good about the performance of this new system for the first year or two, and we had systems in many hospitals around the country. However, we then installed a system in a large, well-known hospital in Chicago, inside "the loop." The equipment was located in a closet on the hospital's third floor, close to the nurses' station. It was not a very comfortable place, but nevertheless, the closet was cozy and had a view of the Chicago skyline.

Minicomputers were then large, and they consumed a lot of power. You had to treat them kindly and pay attention to what kind of environment in which you wanted them to operate. If the minicomputer was unhappy, it would usually just halt. There was no reboot and no watchdog; it just stopped executing code. For this reason, we made sure that the minicomputer rack in our assigned closet had adequate cooling, conditioned power, and other amenities.

The new installation in Chicago would run for days and then randomly stop—apparently without reason. The time, day, date, weather, and phase of the moon could not explain why this would happen. Our service people changed out boards, memory, I/O, cabling, and even the entire minicomputer, but the problem persisted. After a couple of months, we were all baffled, and the hospital employees were getting upset.

I came to work one morning, and my boss strongly suggested that I purchase a one-way ticket to Chicago. He told me not to come back until I had solved this mystery and the hospital was happy. The next day, I found myself in the closet with the minicomputer, where I decided to camp out and see for myself. At least I had a view.

A couple of days went by with no problem. Then it happened! The minicomputer just stopped running, the CRT screens froze, and the keyboard went dead. Nothing seemed to have changed. I restarted the system and waited again. Several hours later, it stopped again. This time, I noticed a small flash in the window. I went to the window in time to notice the "el" train going by. The elevated track ran beside the building at about the level of our floor. After that discovery, I monitored the situation for 24 hours and noticed that the minicomputer stopped only when a train went by.

We demonstrated this problem to the hospital and immediately began grounding and shielding both the equipment and the closet. We installed a wire mesh across the window and the wall and connected it to a newly installed ground that the hospital provided. This approach fixed the problem. Evidently, the train was generating enough EMI (electromagnetic-interference) to stop the minicomputer.

I called my boss and blamed the whole thing on the train. I then asked him whether I could come home.**EDN**

Ray Hill is president of Systemation Technology Inc (Webster, TX), which provides custom embedded-system engineering and consulting, as well as flat-panel graphics interfaces.

Make an impact

Unleash the power of CST STUDIO SUITE®,

the No 1 technology for electromagnetic simulation.

→ Get equipped with leading edge EM technology. CST's tools enable you to characterize, design and optimize electromagnetic devices all before going into the lab or measurement chamber. This can help save substantial costs especially for new or cutting edge products, and also reduce design risk and improve overall performance and profitability.

Involved in signal or power integrity analysis? You can read about how CST technology was used to simulate and optimize a digital multilayer PCB's performance at www.cst.com/pcb. If you're more interested in EMC/EMI or microwave components, we've a wide range of worked application examples live on our website at www.cst.com/apps.

Now even more choice for SI/PI simulation. The extensive range of tools integrated in CST STUDIO

SUITE enables numerous applications to be analyzed without leaving the familiar CST design environment. This complete technology approach enables unprecedented simulation reliability and additional security through cross verification.

→ Grab the latest in simulation technology. Choose the accuracy and speed offered by CST STUDIO SUITE.

CST

CHANGING THE STANDARDS